

Assessing the use of tactical clouds to enhance warfighter effectiveness

Alan Magar
Sphyrna Security Inc.

Prepared By:
Sphyrna Security Inc.
340 Ridgeside Farm Drive
Kanata, ON K2W 0A1

Project Manager: Darcy Simmelink (613) 998-1451
PWGSC Contract Number: W7714-08FE01
Contract Scientific Authority: David Brown (613) 993-9963,
Mazda Salmanian (613) 998-0649
Kathryn Perrett (613) 993-5132

The scientific or technical validity of this Contract Report is entirely the responsibility of the contractor and the contents do not necessarily have the approval or endorsement of Defence Research and Development Canada.

Defence Research and Development Canada

Contract Report
DRDC-RDDC-2014-C69
April 2014

- © Her Majesty the Queen in Right of Canada, as represented by the Minister of National Defence, 2014
- © Sa Majesté la Reine (en droit du Canada), telle que représentée par le ministre de la Défense nationale, 2014

Abstract

Tactical clouds provide a means to extend cloud computing to the tactical environment, in order to effectively move information processing closer to the warfighter. While tactical clouds cannot address the problems (e.g., limited bandwidth, unreliable connections) associated with tactical communications, they do enable warfighters to operate somewhat autonomously by providing a capability to both process data and retrieve historical data locally. They also extend the capabilities of warfighters by providing enhanced functionality not normally available at the edge of tactical networks. This report examined, and assessed, four tactical cloud architectures in order to determine the extent to which cloud computing can effectively be deployed to the tactical edge.

This page intentionally left blank.

Executive summary

Assessing the Use of Tactical Clouds to Enhance Warfighter Effectiveness

Alan Magar; DRDC-RDDC-2014-C69; Defence R&D Canada – Ottawa; April 2014.

Introduction:

The fog war is defined as *the uncertainty in situational awareness experienced by participants in military operations. The term seeks to capture the uncertainty regarding one's own capability, adversary capability, and adversary intent during an engagement, operation, or campaign.*¹ As one would expect, the closer one gets to the tactical edge the more noticeable this phenomenon is. Warfighters at the tactical edge are often forced to make snap decisions in a chaotic environment without sufficient information. This situation is compounded by the fact that communications back to command are over low bandwidth links and are often unreliable.

Tactical clouds provide a means by which a cloud computing capability can be extended to the tactical edge. This capability allows warfighters to access and even process data locally rather than having to rely on constrained communication links back to base. This has the potential to improve warfighter situational awareness and ultimately mitigate the fog of war.

The purpose of this report is to examine the potential for tactical clouds to enhance warfighter effectiveness. Specifically, this was accomplished by examining a number of tactical cloud architectures and assessing their ability to improve overall mission effectiveness.

It should be noted that cloud computing technology, in the form of tactical clouds, is not the only way in which this capability can be provided. Conventional server technology can, and has been, used to provide an information processing capability within the tactical environment. However, the military is increasingly looking at cloud computing technology due to the flexibility that it provides and the ease with which it can be deployed and managed.

Results:

This report identified four tactical cloud architectures that can be used to extend cloud computing into the tactical environment. The four tactical cloud architectures are as follows:

- Centralized Cloud – A centralized cloud, which is usually the result of a data centre consolidation initiative, provides enterprise services to defence personnel. It also allows defence departments to handle the ever-increasing amounts of Intelligence, Surveillance and Reconnaissance (ISR) data generated by sensors. The primary benefits of adopting a centralized cloud architecture are flexibility and cost;

¹ http://en.wikipedia.org/wiki/Fog_of_war

- De-centralized Cloud – A de-centralized cloud, which is typically deployed to a Forward Operating Base (FOB) or onboard ship, is similar in many respects to a centralized cloud except on a significantly smaller scale. It is often a self-contained unit (e.g., shipping container) containing both computing resources and cooling. A de-centralized cloud allows warfighters to process ISR data and analysts to analyze intelligence data locally rather than relying on long distance communication links;
- Cloudlet – Cloudlets, which are “data centres in a box”, are often located in military vehicles in tactical environments. They allow mobile devices carried by warfighters to off-load processing and/or storage. In addition, they can also be used to extend the capabilities of mobile devices. This approach, which is termed Mobile Cloud Computing (MCC), is beneficial in that by offloading these functions to the cloudlet it effectively extends the battery life of the mobile device. Applications that require considerable processing and data access but little bandwidth are ideal candidates for the technology. Vehicular Ad-hoc Networks (VANETs) can be used to interconnect cloudlets and ultimately multiply the processing power and storage available to warfighters; and
- Pico-Cloud – A pico-cloud is a cloud capability that runs on resource constrained devices such as mobile devices. The pico-clouds communicate with one another over a Mobile Ad-hoc Network (MANET) in order to share, process and store data. While the concept of a pico-cloud is interesting in theory, it is unclear how it circumvents the constraints currently imposed by the limited battery life available in mobile devices.

Future plans:

Based on the capability assessments it was determined that cloudlets represent the most viable candidate for further research. Specifically, this report recommended that a detailed examination of this tactical cloud architecture be performed in order to gain a greater understanding of the techniques used for performing computation offloading and capability extension.

Table of contents

Abstract	i
Executive summary	iii
Table of contents	v
List of figures	viii
1 Introduction.....	1
1.1 Background	1
1.2 Purpose	1
1.3 Scope	1
1.4 Assumptions	1
1.5 Disclaimer	1
1.6 Document Structure.....	2
2 Cloud Computing & the Tactical Environment.....	3
2.1 Overview	3
2.2 Tactical Environment	3
2.2.1 Military Concepts	3
2.2.2 Characteristics of the Tactical Environment.....	5
2.2.3 Mobile Device Considerations.....	6
2.3 Tactical Clouds.....	7
2.3.1 Types of Tactical Clouds	8
2.3.2 Tactical Cloud Configurations.....	9
3 Centralized Cloud	11
3.1 Overview	11
3.2 Capability Deficiency	11
3.3 Architectural Overview	12
3.4 Tactical Cloud Discussion.....	13
3.5 Use Cases	14
3.6 Current Initiatives	14
3.6.1 APC2	14
3.6.2 Cloud.mil	15
3.6.3 DISA RACE	15
3.6.4 DISA Storage Cloud	15
3.6.5 DoD Cloud Computing Strategy	16
3.7 Capability Assessment.....	16
4 De-Centralized Cloud	19
4.1 Overview	19
4.2 Capability Deficiency	19
4.3 Architectural Overview	19

4.4	Tactical Cloud Discussion	21
4.5	Use Cases	21
4.6	Current Initiatives	21
4.6.1	APC2	22
4.6.2	DSC	22
4.7	Capability Assessment.....	23
5	Cloudlet.....	24
5.1	Overview	24
5.2	Capability Deficiency	24
5.3	Architectural Overview	25
5.3.1	Terrestrial Cloud	27
5.3.2	Battlegroup Cloud.....	28
5.3.3	Airborne Compute Cloud	29
5.4	Tactical Cloud Discussion	29
5.5	Use Cases	30
5.5.1	MCC Use Cases.....	30
5.5.1.1	Advanced Mapping.....	31
5.5.1.2	Facial Recognition	31
5.5.1.3	Translation Application	32
5.5.1.4	Augmented Reality	33
5.5.2	VANET Use Cases	33
5.5.2.1	Terrestrial Cloud	34
5.5.2.2	Battlegroup Cloud.....	34
5.5.2.3	Airborne Compute Cloud	34
5.6	Current Initiatives	35
5.6.1	MCC Initiatives	35
5.6.1.1	Command and Control Applications for the Decisive Edge	35
5.6.1.2	Connecting Soldiers to Digital Applications (CSDA)	35
5.6.1.3	Enabling Battlefield Decision-Making in the Tactical Cloud	36
5.6.1.4	Joint Battle Command-Platform	36
5.6.1.5	Cloud Computing at the Tactical Edge	36
5.6.1.6	Civilian Research Initiatives	37
5.6.2	VANET Initiatives.....	38
5.6.2.1	VANET and Cloud Research.....	38
5.6.2.2	UAV Research	39
5.7	Capability Assessment.....	39
6	Pico-Cloud	41
6.1	Overview	41
6.2	Capability Deficiency	41
6.3	Architectural Overview	41
6.4	Tactical Cloud Discussion	43

6.5	Use Cases	43
6.6	Current Initiatives	43
6.6.1	Content-Based Mobile Edge Networking	44
6.6.2	Hyrax	44
6.6.3	A Virtual Cloud Computing Provider for Mobile Devices	44
6.7	Capability Assessment	45
7	Future Research	46
8	Conclusion & Recommendations	47
	References	49
	Annex A Cloud Computing Primer	53
A.1	Essential Characteristics	54
A.2	Service Models	55
A.2.1	IaaS	55
A.2.2	PaaS	56
A.2.3	SaaS	56
A.2.4	DaaS	57
A.3	Deployment Models	57
A.3.1	Private Cloud	57
A.3.2	Community Cloud	58
A.3.3	Public Cloud	58
A.3.4	Hybrid Cloud	58
A.4	Key Concepts	59
A.4.1	Abstraction of Resources	59
A.4.2	Multi-Tenancy	61
A.5	The Case for Cloud Computing	61
A.6	Disadvantages of Cloud Computing	63
Annex B	Wireless Communication Primer	65
B.1	Communication Frequencies	65
B.2	Mobile Phone Standards	65
B.3	Wireless Networking Standards	66
	List of symbols/abbreviations/acronyms/initialisms	67

List of figures

Figure 1 – Military Concepts	5
Figure 2 - Types of Tactical Clouds	9
Figure 3 - Tactical Cloud Configurations.....	10
Figure 4 - Centralized Cloud (1 of 2)	12
Figure 5 - Centralized Cloud (2 of 2)	13
Figure 6 - De-centralized Cloud (1 of 2)	20
Figure 7 - De-centralized Cloud (2 of 2)	20
Figure 8 – Cloudlet.....	26
Figure 9 – Mobile Cloud Computing using a Cloudlet	27
Figure 10 - Terrestrial Cloud.....	28
Figure 11 - Battlegroup Cloud.....	28
Figure 12 - Airborne Compute Cloud.....	29
Figure 13 – Pico-Cloud (1 of 2)	42
Figure 14 – Pico-Cloud (2 of 2)	43
Figure 15 - Cloud Computing Concepts	54
Figure 16 - Service Model Control	55
Figure 17 – Virtualization Concepts.....	61

1 Introduction

1.1 Background

Cloud computing has been widely adopted within the private sector as organizations attempt to achieve a business advantage over their competitors. Cloud computing provides a means to achieve a competitive advantage by delivering a more agile Information Technology (IT) infrastructure at a reduced cost. While most government organizations have been somewhat slow in adopting this technology, this is starting to change. This is especially true with respect to defence departments around the world. They are starting to see cloud computing as a means to handle increasing amounts of information and ensure that it is made available in a timely manner to those who require it. This could potentially include warfighters at the tactical edge, who, due to technological advancements, have ever increasing computational resources (e.g., mobile devices) at their disposal.

1.2 Purpose

The purpose of this report is to examine the use of tactical clouds to enhance warfighter effectiveness. Specifically, this will be accomplished by examining a number of tactical cloud architectures and assessing their ability to improve overall mission effectiveness.

1.3 Scope

This report examines tactical cloud architectures, including specific initiatives, and assesses their ability to improve overall mission effectiveness. The report was developed using publically available sources of information accessible over the Internet.

1.4 Assumptions

This report assumes that the reader has a good understanding of cloud computing, including the various service and deployment models. Readers lacking this prerequisite are strongly encouraged to read Annex A (Cloud Computing Primer) prior to reading the rest of the report.

This report also assumes that the reader has a basic understanding of wireless communication, especially as it relates to military environments. Readers lacking a background in wireless communication are encouraged to consult Annex B (Wireless Communication Primer) prior to reading the remainder of the report.

1.5 Disclaimer

Many of the examples of tactical cloud initiatives used throughout this report are U.S. Department of Defense (DoD) initiatives. The intent was not to focus specifically on DoD initiatives. However, this ended up being the case due to the plethora of online references made publicly available.

1.6 Document Structure

This report has the following structure:

- a. Section 1: Introduction, provides an overview of the document;
- b. Section 2: Cloud Computing & the Tactical Environment, examines the tactical environment and tactical clouds;
- c. Section 3: Centralized Cloud, examines centralized clouds including the capability deficiency, architectural overview, tactical cloud discussion, use cases, current initiatives and capability assessment;
- d. Section 4: De-centralized Cloud, examines de-centralized clouds including the capability deficiency, architectural overview, tactical cloud discussion, use cases, current initiatives and capability assessment;
- e. Section 5: Cloudlet, examines cloudlets including the capability deficiency, architectural overview, tactical cloud discussion, use cases, current initiatives and capability assessment;
- f. Section 6: Pico-Cloud, examines pico-clouds including the capability deficiency, architectural overview, tactical cloud discussion, use cases, current initiatives and capability assessment;
- g. Section 7: Future Research, identifies areas of potential future research in terms of tactical clouds;
- h. Section 8: Conclusion & Recommendations, lists the conclusions and recommendations derived from the development of the report;
- i. References, identifies the reference material that was used in the development of this report;
- j. Annex A – Cloud Computing Primer, provides an overview of cloud computing including essential characteristics, service models, deployment models, key concepts, the case for cloud computing, and disadvantages of cloud computing;
- k. Annex B – Wireless Communication Primer, provides an overview of wireless communication concepts of importance to the reader including frequencies, mobile phone standards, and wireless networking standards; and
- l. List of symbols/abbreviations/acronyms/initialisms, provides the long form for all of the acronyms used throughout the report.

2 Cloud Computing & the Tactical Environment

2.1 Overview

This paper will examine a number of architectures that will enable cloud computing technology to be deployed in the tactical environment. However, in order to fully appreciate these discussions, the reader must have a common understanding of a number of terms and concepts. This section of the report will examine the tactical environments and tactical clouds.

2.2 Tactical Environment

The tactical environment is dramatically different from the civilian environment. In order to appreciate the discussions within this report it is important to highlight some characteristics of this environment and agree on some terminology. Specifically, this section of the report will address the following:

- Military Concepts;
- Characteristics of the Tactical Environment; and
- Mobile Device Considerations.

2.2.1 Military Concepts

In order to appreciate the discussions throughout the rest of the report, it is important that the reader have a common understanding of a number of military concepts. These concepts, which are illustrated in Figure 1, include the following:

- Command Headquarters (HQ) – Command HQ is a permanently manned military facility from which operations, including overseas operations, are run;
- Main Operating Base (MOB) – The MOB is a permanently manned, well protected military facility typically located overseas;
- Naval Base (NB) – The naval base is a permanently manned coastal military facility used by ships between missions for docking, maintenance or for restocking purposes;
- Forward Operating Base (FOB) – A FOB is a secured forward military position, often a military base, that is used to support tactical operations. It is supported either by a MOB or directly by the Command HQ;

- Carrier Battle Group (CVBG)² – A CVBG is a naval task force consisting of an aircraft carrier (designated CV) and its escorts. These escorts can include cruisers, warships, destroyers, frigates, corvettes and support ships. It is supported either by a naval base or directly by Command HQ;
- Unmanned Aerial Vehicle (UAV) – A UAV, which is also referred to as a drone, is an aircraft that can either be piloted remotely or operate autonomously. Although originally intended for Intelligence, Surveillance and Reconnaissance (ISR) missions, their role has been expanded to include armed reconnaissance, Suppression of Enemy Air Defences (SEAD), air-to-air combat, and homeland security;
- Light Armoured Vehicle (LAV) – Within this report, a LAV is a generic term used to denote a mechanized infantry vehicle used to support warfighters in ground operations;
- Warfighter – Within this report, the warfighter is a generic term used to refer to dismounted infantry in tactical environments. For the purpose of this report, it is assumed that warfighters will be equipped with mobile communication devices. These are discussed in more detail in Section 2.2.3; and
- Tactical Edge – The *Tactical Edge Characterization Framework – Volume 1: Common Vocabulary for Tactical Environments* [Reference 34] defines the tactical edge from both a user and a technology perspective. *From a user perspective, the tactical edge is defined in the Network Centric Operating Environment (NCOE) Joint Capabilities Document (JCD) as the “First Tactical Mile”. Users are warfighters directly involved in executing the mission at the “tip of the spear”. In this context, the JCD defines “users” as those executing the mission in a forward deployed position. From a technology perspective, the tactical edge is where users operate in certain environments that are constrained by such things as limited communications connectivity and limited storage availability.* Within this report the tactical edge will be used to refer to an imaginary construct separating warfighters from enemy combatants on the field of battle.

² CVGB, which is sometimes referred to as a Carrier Strike Group (CSG), is an American term. There is no equivalent term within the Canadian Navy due to the fact that it does not currently have an aircraft carrier in its fleet.




Figure 1 – Military Concepts ³

2.2.2 Characteristics of the Tactical Environment

The tactical environment is quite different from the modern, civilian environment that many readers are accustomed to. In order to fully appreciate the challenge of extending cloud computing to the tactical edge, it is crucial that the reader understand a number of characteristics of the tactical environment that differentiate it from civilian environments. These characteristics include the following:

- Hostile Environment – The tactical environment is a hostile environment that includes adversaries intent on disrupting communications and harming warfighters. Communication disruptions can be the result of communication equipment that has been destroyed or communication signals that have been jammed. Consequently, warfighters

³ It is worth noting that the various components depicted in Figure 1 are often connected by unreliable and constrained communications networks commonly found in tactical environments. This is discussed in further detail in Section 2.2.2.

- are sometime forced to operate in a “denied environment” (communications detrimentally affected) or an “operating while compromised environment” (attackers with access to communications network);
- Big Data – Since 9/11 the amount of surveillance data generated by drones and other surveillance technology has risen by 1600 percent.⁴ Unfortunately, this rapid increase in sensor data has outpaced the rate at which the data can be transmitted or processed. Consequently, much of the petabytes of data being collected remains unprocessed and therefore unanalyzed. The big data problem is unlikely to improve as it is projected that sensor data volume could potentially be measured in yottabytes (10^{24} bytes) by 2015⁵; and

- Disconnected, Intermittent and Low-bandwidth (DIL) Communications – Modern civilian environments are typically equipped with a considerable amount of fixed communication infrastructure (e.g. cellular towers, telephone lines, wireless access points, fibre optic links) that allows for constant connectivity over relatively high bandwidth networks. In contrast, tactical environments typically lack this fixed communication infrastructure due to either the remote location or the fact that the infrastructure has been rendered inoperative. This lack of a fixed communication infrastructure means that warfighters must rely on potentially unreliable communication links. Even when fully functioning, these links tend to be lower bandwidth than comparative links in modern civilian environments. The challenge of DIL communications becomes increasingly difficult, the closer one gets to the tactical edge.

2.2.3 Mobile Device Considerations

Mobile devices are handheld computing devices that are used by warfighters primarily for communications purposes. Although the majority of mobile devices currently in use are smartphones and Personal Role Radio (PRR), warfighters may eventually be equipped with tablets, wearable devices (e.g., Google Glass) and smart rifles. Due to technological advances, mobile devices are often equipped with more advanced functionality such as digital cameras and Global Positioning System (GPS) navigation systems. They require a communication network, such as a 3G/4G service, WiFi or peer-to-peer networking, in order to communicate directly with other warfighters or back to the FOB. Mobile devices are constrained devices in that they have significantly less computational power and less storage than their desktop equivalents. The following constraints must be considered:

- Size – Mobile devices are limited to a relatively small form factor in order to be portable enough for use by warfighters. These size restrictions have a detrimental effect on performance. Due to their size, mobile devices also have a tendency to be susceptible to loss, theft, and damage;
- Processing – The requirement to make mobile devices smaller and lighter detrimentally affects their computing capabilities. It was determined that servers have 1000x to 5000x

⁴ <http://www.forbes.com/sites/techonomy/2012/03/12/military-intelligence-redefined-big-data-in-the-battlefield/>

⁵ *Data Analysis Challenges* [Reference 5]

more processing power than mobile phones.⁶ Furthermore, even with the enormous technological advancements in mobile technology, the gap between mobile devices and servers has not narrowed. It was determined⁷ that over a 15-year period the gap in the processing power of a typical server and a mobile device was consistent;

- Storage – While the size and cost of flash storage is improving, the amount of storage available on mobile devices is still quite limited. Furthermore, the integrity of data stored on flash storage is diminished by the wear caused by “erase” and “write” operations. This situation is compounded by increasing storage requirements;
- Battery – Battery capacity has failed to increase at the pace of other aspects (e.g., processing, storage) of mobile devices. Worse still, improvements to mobile devices (e.g., faster processors, larger displays, additional functionality (e.g., camera)) actually increase power consumption on mobile devices. It is for these reasons that power constraints are an extremely important consideration in extending cloud computing to the tactical edge, especially in light of the desire to reduce the amount of spare batteries carried by warfighters and the inability to readily recharge batteries.⁸

At any given cost and level of technology, considerations of weight, size, battery life, ergonomics, and heat dissipation exact a severe penalty in computational resources such as processor speed, memory size, and disk capacity. From the viewpoint of a user, a mobile device can never be too small, too light or have too long a battery life. While mobile hardware continues to evolve and improve, it will always be resource-poor relative to static hardware. On hardware that people carry or wear for extended periods of time, improving size, weight and battery life are higher priorities than enhancing compute power. This is not just a temporary limitation of current mobile hardware technology, but is intrinsic to mobility. Computation on mobile devices will always be a compromise.⁹

2.3 Tactical Clouds

A tactical cloud is defined in this report simply as a cloud computing capability that has been made available in a tactical environment. This section will examine the following aspects of tactical clouds:

- Types of Tactical Clouds; and
- Tactical Cloud Configurations.

⁶ *Hyrax: Cloud Computing on Mobile Devices using MapReduce* [Reference 7]

⁷ *The Impact of Mobile Multimedia Applications on Data Centre Consolidation* [Reference 6]

⁸ It was determined that in a typical 72-hour mission in Afghanistan, U.S. soldiers carried seven different types of batteries. These 70 individual batteries (seven different types) added almost 20 pounds to the weight a warfighter is forced to carry in theatre. <http://www.arl.army.mil/www/?article=564>

⁹ *The Impact of Mobile Multimedia Applications on Data Centre Consolidation* [Reference 6]

2.3.1 Types of Tactical Clouds

A centralized cloud provides access to a pool of computing resources typically located in one or more data centres. Since the underlying computing resources are abstracted, dynamic computing requirements can be addressed by simply allocating or de-allocating computing resources. In terms of defence organizations, centralized clouds can be used to provide all of the enterprise computing capabilities typically associated with conventional systems and data centres. The centralized cloud is discussed in Section 3.

If it were possible to access a centralized cloud from anywhere in the world, including from within the tactical environment, there would be little requirement for other types of tactical clouds. However, due to the massive amounts of data involved and the potentially constrained communication links available, centralized clouds are not always readily accessible. Consequently, smaller clouds offering similar services, or a subset of the services, as centralized clouds are required. These tactical clouds, which are illustrated in Figure 2, include the following:

- De-centralized Cloud – A de-centralized cloud is based on the same technology as a centralized cloud but it differs in terms of the scale of the implementation. De-centralized clouds tend to be deployed in geographically remote parts of the world where communication with the centralized cloud is problematic due to inconsistent communication links. De-centralized clouds, including both the computing and cooling functions, are sometimes deployed in shipping containers in order to facilitate deployment. The de-centralized cloud is discussed in Section 4;
- Cloudlet – While a cloudlet is also based on the same technology as a centralized cloud, it is orders of magnitude smaller. A cloudlet is sometimes referred to as a “data centre in a box”. It consists of a number of discoverable, stateless servers running Virtual Machines (VMs). Cloudlet resources, specifically computational and storage resources, are typically leveraged by mobile devices in close geographic proximity in order to off-load resource-intensive computations. The cloudlet is discussed in Section 5; and
- Pico-Cloud – A pico-cloud is the name given to an extremely limited cloud capability provided using resource-constrained devices such as mobile devices. While a pico-cloud is based on the same concepts as a cloudlet, it typically has three orders of magnitude less computational resources available. The pico-cloud is discussed in Section 6.




Figure 2 - Types of Tactical Clouds

2.3.2 Tactical Cloud Configurations

This section of the report will examine two types of tactical cloud configurations; conventional tactical clouds and dynamic tactical clouds. These tactical cloud configurations, which are not mutually exclusive, can be seen in Figure 3.

A conventional tactical cloud uses cloud computing technologies and concepts to provide enhanced computation and storage capabilities to operational environments. It is deemed conventional in that standard computing devices, including mobile devices, are used to access remote clouds. The links between physical computing equipment comprising the cloud are typically high bandwidth. The links from tactical users to the cloud may be low bandwidth.

A dynamic tactical cloud uses cloud computing technologies and concepts to provide enhanced computation and storage capabilities to dynamic operational environments. It is deemed dynamic in that cloud-enabled devices establish ad-hoc networks with one another. Within these ad-hoc networks the number of cloud-enabled devices is constantly changing, as cloud-enabled devices join and depart the dynamic tactical cloud.




Figure 3 - Tactical Cloud Configurations

3 Centralized Cloud

3.1 Overview

This section of the report will examine centralized clouds. In some ways it is a misnomer to include a centralized cloud in a report purporting to examine tactical clouds as strictly speaking, a centralized cloud is primarily strategic in nature. However, the centralized cloud has been included for the sake of completeness and due to the fact that it has a tangible effect on operations. Specifically, this section of the report will examine the following aspects of a centralized cloud:

- Capability Deficiency;
- Architectural Overview;
- Tactical Cloud Discussion;
- Use Cases;
- Current Initiatives; and
- Capability Assessment.

3.2 Capability Deficiency

Defence organizations, perhaps even more so than other modern organizations, are highly dependent on technology, and specifically IT, in order to accomplish their goals. However, the adoption of IT within defence organizations has struggled to keep pace with the dramatic increase in the amount of data requiring processing. This is due in part to the long acquisition times typically associated with defence IT systems. The Defense Science Board ¹⁰ analyzed 32 major automated information system acquisitions within DoD. The average time to deliver an initial program capability was determined to be 91 months once funding was approved. According to the Defense Science Board, this was two to three times the average industry IT refresh cycle time, *making it difficult to keep pace with user needs and technology evolution.*¹¹ In addition, IT comes at a considerable cost. For example, the DoD is forecasted to spend in excess of \$34.5 billion on information technology in 2014.¹² Defence organizations, as with most organizations, are under considerable pressure to reduce IT expenditures.

¹⁰ The Defense Science Board, which was established in 1956, is a committee of civilian experts who provide scientific and technical advice to the DoD.

¹¹ *Cloud Computing Strategy* [Reference 3]

¹² *Government Accountability Office (GAO) – Agencies Need to Strengthen Oversight of Multibillion Dollar Investments in Operations and Maintenance* [Reference 31]

3.3 Architectural Overview

A centralized cloud, which is illustrated in Figure 4 and Figure 5, is a cloud computing capability located back in command HQ, a MOB or a NB. The centralized cloud is considered centralized in that it is usually the result of a data centre consolidation initiative. However, the centralized cloud can still be distributed across a number of data centres in multiple geographic locations. In the distributed configuration, the communication links between cloud computing components are such that there is little discernible difference between configurations. The centralized cloud, which typically provides enterprise services, and is accessible by both local and remote devices. Mobile devices can access the centralized cloud through wireless access points (WiFi) or 3G/4G cellular networks, whereas systems located in the FOB or the CVBG access the centralized cloud over long distance communication links. It is worth noting that these communication links are likely constrained in terms of bandwidth and may be unreliable at times. The centralized cloud will necessitate separate implementations for unclassified, Secret and Top Secret security domains.




Figure 4 - Centralized Cloud (1 of 2)




Figure 5 - Centralized Cloud (2 of 2)

3.4 Tactical Cloud Discussion

Conventional data centres have been used exclusively in the past to provide enterprise services for defence organizations. However, a centralized cloud provides a number of benefits to an organization. These general benefits are discussed in Section A.5. In terms of addressing the capability deficiencies highlighted in Section 3.2, a centralized cloud provides the following specific benefits:

- Big Data – As mentioned, defence organizations have struggled to address the dramatic increase in the amount of data requiring processing. A centralized cloud provides a scalable means with which to process this information. Instead of relying on a finite number of specialized systems, a centralized cloud distributes the processing across a large number of commodity servers. An increase in the amount of data to be processed can simply be addressed by adding additional cloud resources;
- Flexibility – A centralized cloud provides a layer of abstraction between applications and the underlying hardware resources. This allows programs to deploy applications on existing systems and easily scale the program as required; and
- Cost Savings – The implementation of a centralized cloud typically involves the consolidation of a large number of data centres. This endeavour provides a number of opportunities for cost savings. In a traditional data centre most applications typically use a small percentage of a server's processing capacity. In a centralized cloud the utilization rates are in excess of 70%.¹³ Consequently, a centralized cloud will use less hardware than traditional data centres. Less hardware translates directly to savings in terms of power utilization, space requirements, acquisition costs and support costs.

¹³ According to Gartner (*Energy Savings via Virtualization: Green IT on a Budget* [Reference 33]) most x86 servers run at performance levels of between 7% and 15%, while target performance levels for virtualization should be in the 65 to 70% range.

3.5 Use Cases

The following use case has been included in order to illustrate the concept of a centralized cloud and describe how the technology can make a difference in the tactical environment:

A group of warfighters will be traveling into hostile territory, which is occupied by insurgents, in order to find caches of weapons and confiscate them. However, prior to entering into hostile territory the warfighters require specific data on likely locations for the weapons caches. Furthermore, the warfighters are concerned about the possibility of an armed ambush or concealed Improvised Explosive Device (IED). Consequently, they require guidance in terms of likely locations for armed ambushes and IEDs.

Analysts perform a search on the centralized cloud database of ISR information, both current and historical, for insurgent activity in the area. Although the ISR database comprises petabytes of data, analysts are able to perform their searches relatively quickly. In a matter of a few hours the analysts have identified a number of potential locations for weapons caches. Furthermore, the analysts were able to pinpoint a number of ambush and IED locations. This information is shared with the warfighters prior to their deployment from the FOB, thereby increasing the warfighters chances of successfully completing their mission

3.6 Current Initiatives

While there are countless civilian centralized cloud initiatives, there are significantly less within the military. Current centralized cloud initiatives include the following:

- Army Private Cloud (APC2);
- Cloud.mil;
- Defense Information Systems Agency (DISA) Rapid Access Computing Environment (RACE);
- DISA Storage Cloud; and
- DoD Cloud Computing Strategy.

3.6.1 APC2

APC2, which is a component of the LandWarNet strategic initiative, is a five year, \$249.8 million dollar fixed-price contract awarded to IBM, Lockheed Martin, HP, General Dynamics, Northrop Grumman, MicroTech and Criterion Systems. The contract consists of two parts. The first part, which will be discussed in this section of the report, is a data centre consolidation project that will result in a private, centralized cloud. The second part, which will be discussed in Section 4.6.1, will result in mobile, containerized data centres suitable for deployment in FOBs.

The data centre consolidation portion of APC2 is intended to consolidate data centres from more than 200 to fewer than 20. The U.S. Army believes that the centralized cloud will result in lower application migration, hosting, administration and maintenance costs through the provision of on-

demand IT services and storage (i.e., Infrastructure-as-a-Service (IaaS)). However, critics are less optimistic given that it only amounts to 0.6% of the Army's IT budget (average of \$50 million per year divided by a 2009 IT budget of \$7.8 billion). The implication being that the amount budgeted is insufficient to properly implement the centralized cloud that is envisioned.

3.6.2 Cloud.mil

Cloud.Mil is a private Platform-as-a-Service (PaaS) solution developed by the DoD Software Engineering Center (SEC) that is intended to streamline the application development process within DoD. Normally, application development requires hardware, an operating system, databases, middleware, Web servers, and software development environments (e.g., J2EE, .NET). The development team needs to address all aspects of application development including infrastructure licensing, scalability, security, testing, accreditation, etc. Cloud.Mil facilitates this process by providing a self-service, on-demand development and deployment platform for developers. Not only does Cloud.Mil expedite the application development process but it streamlines the Information Assurance (IA) accreditation process while increasing the overall security posture of the applications. In addition, Cloud.Mil provides an administrative capability that performs basic functions, including backup and recovery, log monitoring and performance metrics. The end result is that application developers can focus on the development of their application rather than all of the other aspects associated with application development.

3.6.3 DISA RACE

DISA¹⁴ RACE is an IaaS cloud computing and virtualization infrastructure for developing, testing, and deploying new DoD applications. The test and development cloud, which was developed by HP and has been operating since 1 October 2008, has been used to develop and test hundreds of military applications. These applications include command and control systems, convoy control systems, and satellite programs. Through the use of an online portal DoD customers can purchase computer operating systems, applications and services, and pay for them with a government credit card. With DISA RACE the acquisition time for a new server has been reduced from six months to 24 hours. This is due to the fact that application developers do not need to acquire their own servers and install the requisite software. Instead, they merely use a subset of the DISA RACE resources that are automatically provisioned for their use.

3.6.4 DISA Storage Cloud

The DISA Storage Cloud project is a \$45 million project for a private storage cloud for intelligence and surveillance imagery. The resulting Large Data Object (LDOS) cloud service will allow the agency to securely store hundreds of billions of objects in a way that users could access the data across multiple networks. Data being stored in the cloud would consist of hundreds of billions of imagery files, including standard and high-definition (HD) video, lidar images, infrared and electro-optical images, Wide-Area Motion Imagery (WAMI), and Full Motion Video (FMV). While the initial contract is for 10 petabyte (Pbyte) units tied together via an Internet Protocol (IP) network and hosted in a secure data centre, it is believed that it could eventually provide four exabytes of storage. The contract was originally awarded (sole-source) to Alliance

¹⁴ DISA is a combat support agency of DoD.

Technology Group in April 2013. However, that contract was cancelled in May 2013 in order to allow it to be competed.

3.6.5 DoD Cloud Computing Strategy

The U.S. DoD Cloud Computing Strategy, which was released in July 2012, *introduces an approach to move the Department from the current state of a duplicative, cumbersome, and costly set of application silos to an end state which is an agile, secure, and cost effective service environment that can rapidly respond to changing mission needs.*¹⁵ The approach to transition the department to a DoD-wide centralized cloud consists of the following four steps:

- 1) Foster Adoption of Cloud Computing – *by establishing a strong governance structure that has the authority and responsibility to drive an Enterprise-First approach and enable IT financial, acquisition, and contracting policy and practice reforms;*
- 2) Optimize Data Centre Consolidation – *by implementing a limited set of standardized software platforms and data centers that will enable effective management as a single enterprise with a reduced intrusion surface for cyber threats;*
- 3) Establish the DoD Enterprise Cloud Infrastructure – *as the foundation for rapid participation in the DoD Enterprise Cloud Environment; and*
- 4) Deliver Cloud Services – *using commercial service providers and continuing the development and implementation of DoD cloud services.*

The DoD Cloud Computing Strategy specifies separate implementations and data exchanges for Non-secure Internet Protocol Router Network (NIPRNet), Secure Internet Protocol Router Network (SIPRNet), and Top Secret Sensitive Compartmentalized Information (TS SCI) security domains.

3.7 Capability Assessment

*The underlying value proposition of cloud computing is that centralization exploits economies of scale to lower the marginal cost of system administration and operations.*¹⁶

There is little doubt that a centralized cloud offers enormous potential in terms of cost savings, flexibility and the ability to process the big data associated with ISR. One has only to look at the widespread adoption of this technology within the private sector in order to appreciate its potential within the military. In terms of the tactical environment, a centralized cloud provides a number of secondary benefits. It would help to facilitate secure information sharing and collaboration, and as a result enhance mission effectiveness. This would be accomplished by facilitating the rapid development and deployment of improved applications that could be leveraged in the tactical environment. In addition, the parallel processing provided by a centralized cloud will facilitate advanced analytics on large data sets, thus providing warfighters with the requisite information if a link exists to this cloud. Lastly, the data platform provided by a

¹⁵ DoD Cloud Computing Strategy [Reference 3]

¹⁶ The Impact of Mobile Multimedia Applications on Data Centre Consolidation [Reference 6]

centralized cloud is not only resilient but capable of distributing data across various mission environments.

In terms of the maturity evaluation factor¹⁷, the centralized cloud scores a 5 since the technology has been commercially available for a number of years. In terms of the current level of adoption evaluation factor, the centralized cloud scores a 3 since the technology has some deployment within the military environment (not specifically adopted in a tactical environment). In terms of the current impact evaluation factor, the centralized cloud scores a 2 since the technology has had little impact on warfighter effectiveness thus far. In terms of the potential impact evaluation factor, the centralized cloud scores a 3 since the technology will likely have some impact on warfighter effectiveness in the future. It is important to note that both the effectiveness and the potential of the centralized cloud suffer somewhat due to the fact that it is being assessed on its impact on warfighter effectiveness. While a centralized cloud will in all likelihood have a dramatic impact on IT operations, this does not necessarily guarantee that the impact makes it all the way to the tactical edge. The capability assessment for the centralized cloud can be seen in Table 1.

Table 1 - Centralized Cloud Capability Assessment

Evaluation Factors	Centralized Clouds
Maturity	5
Current Level of Adoption	3
Current Impact	2
Potential Impact	3

Note – Capability Assessment Methodology

The tactical cloud architectures will be assessed according to four evaluation factors that are each meant to stand on their own. Up to five points will be assigned for each evaluation factor. The four evaluation factors consist of maturity, current level of adoption, current impact and potential impact.

The maturity evaluation factor is used to assess the overall level of maturity of the technology. This evaluation factor encompasses both civilian and military technological advances. The maturity ratings are as follows:

- 0 - A score of zero denotes that the technology is theoretical only;
- 1 - A score of one denotes that the technology is currently undergoing limited research and prototyping;
- 2 - A score of two denotes that the technology is undergoing extensive research and prototyping;
- 3 - A score of three denotes that the technology has been commercialized;
- 4 - A score of four denotes that the technology is being offered by a number of vendors; and

¹⁷ For a discussion of the interpretation/meaning of the evaluation factors, please consult the note included at the end of this section.

- 5 - A score of five denotes that the technology has been commercially available for a number of years.

The current level of adoption evaluation factor is used to assess the overall level of adoption of the technology in the tactical environment. The adoption ratings are as follows:

- 0 - A score of zero denotes that the technology has not been deployed in the tactical environment;
- 1 - A score of one denotes that the technology has extremely limited deployment in the tactical environment;
- 2 - A score of two denotes that the technology has limited deployment in the tactical environment;
- 3 - A score of three denotes that the technology has some deployment in tactical environment;
- 4 - A score of four denotes that the technology has widespread deployment in the tactical environment; and
- 5 - A core of five denotes that the technology is pervasive throughout the tactical environment.

The current impact evaluation factor is used to determine the current impact of the technology on warfighter effectiveness. The effectiveness ratings are as follows:

- 0 - A score of zero denotes that the technology has had a negligible impact on warfighter effectiveness;
- 1 - A score of one denotes that the technology has had a minimal impact on warfighter effectiveness;
- 2 - A score of two denotes that the technology has had little impact on warfighter effectiveness;
- 3 - A score of three denotes that the technology has had some impact on warfighter effectiveness;
- 4 - A score of four denotes that the technology has had significant impact on warfighter effectiveness; and
- 5 - A score of five denotes that the technology has had a profound impact on warfighter effectiveness.

The potential impact factor is used to determine the potential impact of the technology on warfighter effectiveness. The potential ratings are as follows:

- 0 - A score of zero denotes that the technology will likely have a negligible impact on warfighter effectiveness;
- 1 - A score of one denotes that the technology will likely have a minimal impact on warfighter effectiveness;
- 2 - A score of two denotes that the technology will likely have little impact on warfighter effectiveness;
- 3 - A score of three denotes that the technology will likely have some impact on warfighter effectiveness;
- 4 - A score of four denotes that the technology will likely have a significant impact on warfighter effectiveness; and
- 5 - A score of five denotes that the technology will likely have a profound impact on warfighter effectiveness.

4 De-Centralized Cloud

4.1 Overview

This section of the report will examine de-centralized clouds. Specifically, this section of the report will examine the following aspects of de-centralized clouds:

- Capability Deficiency;
- Architectural Overview;
- Tactical Cloud Discussion;
- Use Cases;
- Current Initiatives; and
- Capability Assessment.

4.2 Capability Deficiency

Dependable communications between a FOB or CVBG and the main base or HQ to which they are reporting are extremely important for a variety of reasons. In many cases information obtained in the operational environment must be sent back to the main base for either processing or so that decisions can be made based upon the information. In the case of processing, the raw data can be extremely large if it is obtained from high-resolution sensors. Likewise, information is often sent from the main base to the FOB or CVBG. This information could include updated directives or mission data. However, in many cases the communication link connecting these operational entities with geographically distant bases are limited in terms of bandwidth and are often unreliable. This could result in a delay in receiving information, including mission data, required to successfully execute the mission.

4.3 Architectural Overview

As illustrated in Figure 6 and Figure 7, a de-centralized cloud can be deployed in a FOB or onboard a ship. Systems and devices in close proximity (in the FOB or onboard ship) would access the de-centralized cloud directly. In addition, systems located in a LAV or even a UAV might also be able to access cloud resources located back in the FOB. A de-centralized cloud should be easy to deploy. For example, the computing and cooling elements could be included in a single cargo container.




Figure 6 - De-centralized Cloud (1 of 2)




Figure 7 - De-centralized Cloud (2 of 2)

4.4 Tactical Cloud Discussion

While a de-centralized cloud does nothing to resolve the uncertainty of long distance communication links, it can help to mitigate their effect by enabling local processing, information access and decision-making. Although these functions could also be accomplished using conventional servers, deployment and maintenance of these resources can be costly and problematic. Specifically, the de-centralized cloud will provide the following benefits:

- Local Processing - Big data, consisting of intelligence imagery and video, can be processed and analyzed locally rather than having to be sent back to a main base. This local processing will ultimately improve response time and allow actions to be taken sooner;
- Information Access - Troops can be deployed with relevant historical intelligence data, including imagery, required to complete their mission; and
- Decision-Making – By deploying analysts along with the cloud computing capability it allows commanders to leverage local analysts. While the chain of command would still apply, it would enable commanders to make some decisions locally rather than always having to defer to their superiors back at the main base.

4.5 Use Cases

The following use case has been included in order to illustrate the concept of a de-centralized cloud and describe how the technology can be used in the tactical environment:

The Army is establishing a FOB near hostile territory in a remote location of the globe. Communications back to command HQ are unreliable. Consequently, there is a significant delay in terms of receiving updated intelligence. Even when the link is up, the transfer rate is extremely poor. Fortunately, the commander insisted that a de-centralized cloud be deployed with them. The de-centralized cloud includes all intelligence records and source feeds for the region for the past decade. It also allows them to download and process all ISR data from the UAVs and ground sensors that have already been deployed. The end result is that analysts can perform on-site analysis and provide warfighters with the information they require without having to depend on an unreliable link back to the centralized cloud at command HQ.

4.6 Current Initiatives

Current de-centralized cloud initiatives include the following:

- APC2; and
- Distributed Common Ground System – Army (DCGS-A) Standard Cloud (DSC).

4.6.1 APC2

APC2, which was originally discussed in Section 3.6.1, consists of two parts. The second part is a containerized data centre suitable for deployment. One of the vendors working on the contract sells a Performance Optimized Datacentre (POD). They are preconfigured with racks, cabling and equipment for power and cooling. Depending on the model they can be shipped in six to twelve weeks.

4.6.2 DSC

DCGS is DoD's primary system for posting, processing and disseminating intelligence, surveillance and reconnaissance information. While each service has its own version, they are all built on a common architecture in order to facilitate information sharing. It contains data from 409 different sources, including space, airborne and terrestrial assets. This includes a database with detailed intelligence on the threat posed by IEDs, as well as biometrics data collected by individual soldiers. All told, the system contains every intelligence report since 2003, a full petabyte of data. DCGS-A was originally deployed in Afghanistan in 2006.

The Army dramatically increased the robustness of the system five years later (Spring 2011) by adding a cloud computing capability. DCS is capable of consuming 10 gigabytes of raw data per hour and processes queries in an average of 1.3 seconds. It indexes and stores text and visual information on upwards of 75 million intelligence records from as many as 600 source feeds, such as UAV, satellite imagery and ground sensors. This includes every written intelligence report filed by U.S. forces in Afghanistan since Operation Enduring Freedom began in 2001. DSC fits in one cargo container that can be transported on a C-141 Starlifter. DSC provides a number of benefits, including the following:

- Multi-dimensional Analysis of Data – DSC is capable of analyzing and fusing data from multiple sources, including human intelligence, in order to get an accurate picture of the situation;
- Ubiquitous Access – The cloud scales in such a way that it is accessible to anyone with connectivity who needs to gain access to it. This includes those with mobile devices, which access it through a Web browser. Furthermore, by putting the same capability in a smaller form factor, it can be deployed to remote forward bases; and
- Sufficient Resources – Since there are sufficient resources, including processing and storage, analysts no longer have to bracket data. They can conduct real-time analysis of all of the pertinent data.

The Army deployed the first tactical “decentralized cloud” node to Bagram Airfield in northern Afghanistan in mid 2011. This cloud is available to any U.S. or coalition warfighter with a Web browser and with access to either SIPRNet or the coalition network. Another node was deployed in Kandahar in the south. The first node supports Regional Command – East while the second supports Regional Command – South. DSC nodes connect to regional DSC hubs and other DSC nodes in order to increase access to relevant data.

4.7 Capability Assessment

A de-centralized cloud offers significant benefit to remotely deployed warfighters. In terms of deployment, it can be deployed in a matter of weeks instead of months or even years. Once deployed, it allows analysts to store and analyze intelligence data locally rather than relying on long distance communication links. It also allows data to be processed closer to the sensors and the actual conflict. In terms of processing, the de-centralized cloud supports the parallel processing of big data. This allows big data to be sub-divided, processed on cloud resources, and have the answer reassembled once the processing is complete. All of these benefits ultimately ensure that actionable intelligence ends up in the hands of the warfighter quicker than otherwise would have been the case if forced to rely on the centralized cloud and the associated intermittent communication links.

In terms of the maturity evaluation factor, the de-centralized cloud scores a 4 since the technology is being offered by a number of vendors. In terms of the current level of adoption evaluation factor, the de-centralized cloud scores a 2 since the technology has limited deployment within the tactical environment. In terms of the current impact evaluation factor, the de-centralized cloud scores a 3 since the technology has had some impact on warfighter effectiveness (e.g., DSC in Afghanistan). In terms of the potential impact evaluation factor, the de-centralized cloud scores a 4 since the technology will likely have significant impact on warfighter effectiveness. The capability assessment for the de-centralized cloud can be seen in Table 2.

Table 2 – De-centralized Cloud Capability Assessment

Evaluation Factors	De-centralized Clouds
Maturity	4
Current Level of Adoption	2
Current Impact	3
Potential Impact	4

5 Cloudlet

5.1 Overview

This section of the report will examine cloudlets. Specifically, this section of the report will examine the following aspects of cloudlets:

- Capability Deficiency;
- Architectural Overview;
- Tactical Cloud Discussion;
- Use Cases;
- Current Initiatives; and
- Capability Assessment.

5.2 Capability Deficiency

In order to communicate back to base, motorized and mechanized infantry are forced to rely on relatively low-bandwidth communication links that sometimes prove unreliable due to environmental conditions (e.g., mountainous terrain, tall buildings) or enemy activity (e.g., jamming). The situation is even more acute for infantry squads, as an eight-member light infantry section is typically equipped with only one radio capable of communicating back to base.¹⁸ This can have a detrimental effect on situational awareness for both the warfighter and commanders in the FOB. Since the warfighter is sometimes unable to receive situation reports (sitreps) providing them with an update as to the current military situation, they are forced to act without the benefit of this information, potentially affecting chances of mission success. Similarly, data that is collected at the tactical edge must be sent back to the FOB or even the HQ for processing. The elapsed time between collection and processing can be the difference between mission success and failure. Given the constrained communications links and the sheer size of the data being collected, it is infeasible to send all of the data back to a centralized facility for processing.

In many cases, mission data (e.g., maps) can be pre-loaded on warfighter mobile devices prior to departure from the FOB. Due to the constrained nature of mobile devices, a finite amount of information can be stored. If the mission deviates from the original plan then the mission data will need to be updated to reflect the new reality. However, receiving this updated information, including high-resolution images, over low bandwidth, intermittent communication links is a challenge. Furthermore, mobile devices commonly used by warfighters are ill equipped to handle new, resource-intensive applications. Specifically, the processing, storage, and especially the energy limits are often insufficient. The overall impact is that since the information that is required at the tactical edge is unavailable it will have a detrimental impact on mission success.

¹⁸ It is worth mentioning that infantry members are currently equipped with PRR, and that the future vision is to equip them with an even better radio in order to facilitate intra-section communications.

5.3 Architectural Overview

A cloudlet is a data centre in a box. Cloudlets are discoverable, stateless servers running VMs that are in relatively close proximity to the mobile devices whom they are intended to serve. Mobile devices leverage the computational and storage resources of the cloudlet, which are on the order of a thousand times higher than the mobile devices, in order to offload resource-intensive computations. In other words, the mobile device functions as a thin client with the vast majority of the computation being offloaded to the nearby cloudlet. Cloudlets, which are illustrated in Figure 8, are stateless in that they typically require provisioning by a mobile device prior to being able to provide a service. Furthermore, once the cloudlet has been provisioned it does not need to have regular communication with the larger cloud. It can work in a disconnected mode. Close proximity means that the cloudlet is usually a single hop away from the mobile device and can be accessed over WiFi or short-range radio. It is essential that cloudlets be in close physical proximity in order to ensure fast and predictable end-to-end response time of applications executing in the cloudlet. If no cloudlet is available or a cloudlet is lost then the mobile device needs to fall back to its own resources. From a military perspective, cloudlets located in LAVs would enable warfighters to process data nearby while conserving battery power on their constrained mobile devices. This architecture, whereby mobile devices leverage the resources of nearby cloudlets, is sometimes referred to as Mobile Cloud Computing (MCC). MCC, which is illustrated in Figure 9, is a model for elastic augmentation of mobile device capabilities via ubiquitous wireless access to cloud storage and computing resources, with context-aware dynamic adaption to changes in the operating environment.¹⁹ It is a variation on cyber foraging (see note below). However, instead of merely leveraging servers it leverages cloud computing resources, and specifically a cloudlet.

While a cloudlet can provide certain capabilities on its own, its capabilities can be amplified by combining multiple cloudlets in a network. In this scenario, a Vehicular Ad-hoc Network (VANET (see note below)) is used to network cloud-equipped vehicles.²⁰ The network is ad-hoc in that vehicles will join and exit the network as they come in and out of range. As one would expect, the computing and/or storage ability of the networked cloudlets will increase directly with the number of participating nodes. For example, a processing task may take half the time in a VANET with eight nodes as compared to one with four nodes. This report will examine three variations of VANETs consisting of cloud-equipped vehicles. These include the following:

- Terrestrial Cloud;
- Battlegroup Cloud; and
- Airborne Compute Cloud.

¹⁹ *Mobile Cloud Computing: A Comparison of Application Models* [Reference 9]

²⁰ A cloud equipped vehicle is a term used to refer to a vehicle equipped with a cloudlet.

Note – Cyber Foraging

The term cyber foraging, which originated in the paper *Pervasive Computing: Vision and Challenges* [Reference 11], refers to the practice of augmenting the capabilities of resource-constrained mobile devices by offloading resource-intensive tasks to a wired hardware infrastructure such as a conventional data centre. The theory being that through cyber foraging one can extend the battery life of the mobile device, increase computational capability or even provide access to external resources. However, cyber foraging does not address the challenges of operating in dynamic environments that are characterized by unreliable networks.

Note – VANET

A VANET is an ad-hoc network comprised of vehicles. While VANETs share a number of properties in common with Mobile Ad-hoc Networks (MANETs), VANETs have slightly different characteristics due to node mobility (faster) and topological transformations resulting from the high velocities involved. Furthermore, the cloudlet constituents found in VANETs are considerably more powerful than the pico-clouds (Section 6) found in MANETs.




Figure 8 – Cloudlet




Figure 9 – Mobile Cloud Computing using a Cloudlet

5.3.1 Terrestrial Cloud

A terrestrial cloud, illustrated in Figure 10, is an ad-hoc network comprised of cloudlets located in military vehicles (e.g., LAVs). The cloudlets in each vehicle can communicate with one another through the Vehicle-to-Vehicle (V2V) network or even to stationary communication units through Vehicle-to-Infrastructure (V2I) communications. In a military context, a roadside communications unit might provide C2 instructions, serve to upload vehicle maintenance and health status, or even update mapping data. The cloudlet resources are shared via the VANET in order to provide enhanced computation and storage capabilities. In the civilian world, VANETs leverage cellular networks. In fact, the U.S. Federal Communications Commission (FCC) has allocated 75 MHz of spectrum in the 5.850 to 5.925 GHz band for dedicated short-range communications for Intelligent Transportation Systems (ITS).




Figure 10 - Terrestrial Cloud

5.3.2 Battlegroup Cloud

A battlegroup cloud, illustrated in Figure 11, is an ad-hoc network comprised of cloud resources (may be cloudlets or even de-centralized clouds) located in ships that are part of the same battle group. The cloud resources in each ship can communicate with one another through the V2V network in order to provide enhanced computation and storage capabilities.




Figure 11 - Battlegroup Cloud

5.3.3 Airborne Compute Cloud

An airborne compute cloud, illustrated in Figure 12, is an ad-hoc network comprised of cloud resources, basically a micro-data centre, located in an aircraft (e.g., UAV). The cloudlets in each UAV can communicate with one another through the V2V network or even to ground units through V2I communications. The cloudlet resources are shared via the VANET in order to provide enhanced computation and storage capabilities. The UAVs communicate with ground sensors in order to collect ground sensor data. They also communicate with FOBs in order to download the processed sensor data. In both cases, the communications link is a bandwidth-constrained wireless network with variable, unpredictable link qualities. An airborne compute cloud differs from a terrestrial cloud in that the resulting topologies are much more varied. The topologies of a terrestrial cloud have a tendency to conform to the topology of available roads. However, this is not always the case with military terrestrial vehicles. In addition, airborne environments are considered to be limited in terms of size, weight and power when compared to other VANET cloud variations.




Figure 12 - Airborne Compute Cloud

5.4 Tactical Cloud Discussion

While a cloudlet does not resolve all of the capability deficiencies identified in Section 5.2, including potentially unreliable communication links that have a detrimental effect on situational awareness, they do provide the ability to process information, including ISR data, locally. Although this will do little to improve situational awareness at the FOB, it will allow warfighters to better operate autonomously. Not only can warfighters leverage a cloudlet in order to process data but they can retrieve existing data pre-loaded on the cloudlet. In addition, cloudlets interconnected using a VANET provide the ability to distribute data processing/storage amongst

the various participating nodes. Furthermore, by participating in the VANET, nodes extend the range to encompass additional nodes.

A cloudlet provides a cloud computing infrastructure that can be used for a variety of processing tasks. However, in some instances there may be a requirement for specific processing as opposed to general processing. For example, a Mobile Electronic Warfare Team (MEWT) is commonly equipped with servers designed and configured for this specific task. In this case, conventional servers may be a better fit than a cloudlet. As a general rule of thumb, if you know exactly what you want to do, do not need to scale, and there are no other users or applications, then conventional servers are likely better suited to the task than a tactical cloud. However, if you have multiple applications and require additional processing then cloudlets interconnected using a VANET may be preferable due to the ease with which the ad-hoc network can be established and the processing can be distributed amongst participating nodes.

Note – In order to effectively distribute processing and storage between nodes, there will need to be a certain level of connectivity between cloudlets. If there is insufficient connectivity, then the node will be dropped from the VANET and processing/storage will be distributed amongst the remaining nodes.

5.5 Use Cases

In order to best illustrate the complete range of uses, the cloudlet use cases have been divided into two types – those focusing on MCC and those focusing on VANETs.

5.5.1 MCC Use Cases

For MCC, there are a few different use cases that could be examined. While ultimately four were selected, other possible use cases include a direction finding fusion centre, image recognition, mission planning, route calculation, injury diagnosis, etc. The four use cases to be examined are as follows:

- Advanced Mapping;
- Facial Recognition;
- Translation Application; and
- Augmented Reality.

Note – It is important to note that the use cases included in this section of the report describe specific processing tasks that could be addressed using conventional servers. However, cloudlets provide the ability to support multiple use cases rather than just providing a single capability.

5.5.1.1 Advanced Mapping

The following use case has been included in order to illustrate the concept of a cloudlet facilitating advanced mapping and describe how the technology can be used in the tactical environment:

While patrolling, a group of warfighters receives a call for help from another patrol who are patrolling a neighbouring sector. Apparently, the neighbouring patrol is in a firefight with some insurgents. They require aid because they were ambushed and are significantly outnumbered. The warfighters quickly hurry to the sector of the ambushed patrol, aided by their mobile devices that are in direct communication with a cloudlet. Their mobile devices display a detailed map with the user's current location and customized directions. Upon nearing the conflict the map changes to a detailed view of the battlefield with friendlies in blue and insurgents in red. The warfighters can switch to a live video feed from a UAV circling overhead. In addition, by clicking on any of the blue dots, the mobile device will display video from the wearable mobile device and/or scope of the smart rifle of that particular individual. With the advanced mapping provided on his mobile device, the sergeant leading the warfighters is able to deploy his warfighters in such a way as to quickly rout the insurgents.

Note – In this scenario the cloudlet is used for a number of tasks. First and foremost, the cloudlet receives the coordinates from each mobile device carried by the warfighters and superimposes them over a detailed map of the area. It should be noted that there are other ways to accomplish this aspect of advanced mapping including using peer-to-peer broadcasts. In addition, the cloudlet is able to incorporate feeds from other sources, including the UAV, and process this information locally. The processed results can then be displayed on the warfighters' respective devices. There are a number of potential applications. For example, the cloudlet could identify insurgents and determine their location by comparing the UAV feed with the known location of the warfighters. Any combatant in a location other than those identified by the warfighters' mobile devices would automatically be identified. The cloudlet could also take the live feed from the UAV and have it sent directly to warfighters' devices in order to provide a bird's eye view of enemy positions.

5.5.1.2 Facial Recognition

The following use case has been included in order to illustrate the concept of a cloudlet facilitating facial recognition and describe how the technology can be used in the tactical environment:

A group of warfighters are on patrol near an unfamiliar village. Prior to entering the village they send a query to the MCC asking it for the identity of the village elders. Consequently, prior to entering the village they have pictures to help them recognize the village elders and additional information about the elders, including previous contact information and whether any of them speak English.

While on patrol in the village, their wearable mobile devices automatically take pictures of the villagers that they encounter and send them to the cloudlet for processing. The cloudlet attempts to identify the villagers using facial recognition and its extensive database of images. One of the villagers is identified as a known insurgent with instructions to apprehend immediately. The warfighters take the villager into custody.

Note – According to *Cloud Offload in Hostile Environments* [Reference 12] facial recognition played a key role in helping Navy SEALs to positively identify Osama bin Laden prior to undertaking the mission in Abbottabad, Pakistan.

5.5.1.3 Translation Application

The following use case has been included in order to illustrate the concept of a cloudlet facilitating voice translation and describe how the technology can be used in the tactical environment:

While in the village the group of warfighters would like to talk with the village elders in order to gather information about insurgents in the area. Unfortunately, none of the villagers speaks any English or Arabic. Using the MCC a soldier can speak English into their mobile device and then hit a button to have it translated into Pashto or Dari. Likewise, the villagers can speak into the mobile device, and MCC will translate it back to English for the soldier. Using this translation application the warfighters are able to have a productive meeting with the village elders resulting in some leads as to where to find the insurgents.

Note – It should be noted that any language spoken into the mobile device is recorded, sent to the MCC, translated by the MCC, and then sent back to the mobile device. The actual translation is performed by the cloudlet due to the processing intensive nature of the task. Mobile devices would be unable to perform this task locally due to the processing limitations of the devices themselves.

Note – TRANSTAC

At the behest of the Defense Advanced Research Projects Agency (DARPA), a team of scientists from NIST designed a translation system dubbed TRANSTAC (spoken language communication and TRANSlation system for TACTical use). TRANSTAC was a five year program that utilized a smartphone for real-time translation. The software, which supported Pashto, Arabic and Dario, achieved an 80% accuracy rate. While it was used by a few dozen users in Iraq and Afghanistan, *no one was impressed enough to want to keep it.*²¹ Ultimately TRANSTAC was deemed interesting but not useful due to its 80% accuracy rate.

TRANSTAC has been developed into a product by SRI International. The translation system, called IraqComm, consists of three primary software components: Automatic Speech Recognition (ASR), Machine Translation (MT), and Text-To-Speech synthesis (TTS). The ASR model takes voice audio and converts it into text. The MT translates the text into the target language. The TTS then “reads” the translated text. IraqComm is capable of running on a desktop, notebook, tablet and handheld PC.

While TRANSTAC and IraqComm are an example of a translation application, they did not use MCC. It is worth mentioning that the use of MCC could in all likelihood significantly improve the accuracy rate, and therefore its overall usefulness to warfighters.

²¹

http://www.slate.com/articles/technology/future_tense/2012/05/darpa_s_transtac_bolt_and_other_machine_translation_programs_search_for_meaning_.html

5.5.1.4 Augmented Reality²²

The following use case has been included in order to illustrate the concept of a cloudlet providing an augmented reality capability and describe how the technology can be used in the tactical environment:

A patrol is entering a village known to harbour insurgents. As the warfighters walk through the village, street names and building identification information are superimposed on the image that the soldiers see through their wearable devices, helping them to quickly navigate. As the sun sets the wearable devices display thermal imaging information, helping soldiers to identify all heat sources, including hidden insurgents. Another patrol, which has recently entered the village, appears with a blue outline around them to provide a visual indicator to the warfighters that these are friendly forces.

Note – In this scenario the cloudlet processes the data retrieved by the wearable devices. Specifically, it takes the data, including thermal imaging data, and superimposes building identification information based upon the warfighter's location. It is also able to differentiate between friendly and potentially hostile forces, and highlight this difference, as the locations of friendly forces are known.

Note – Cloud-Mobile Convergence (CMC)

The term CMC, which originated in the paper *The Cloud-Mobile Convergence Paradigm for Augmented Reality* [Reference 13], is based on the observation that present day mobile devices are surrounded by abundantly available cloud resources, and access to these resources has never been as convenient before. This definition highlights the difference in providing advanced mobile-cloud capabilities, such as AR, in modern society as compared to tactical environments.

5.5.2 VANET Use Cases

In order to illustrate the potential of ad-hoc networks of cloudlets, it was decided to include a use case for each of the three variations of cloudlet VANETs. The three variations are as follows:

- Terrestrial Cloud;
- Battlegroup Cloud; and
- Airborne Compute Cloud.

²² Augmented Reality (AR) superimposes computer generated graphics over a live image in order to provide additional information.

5.5.2.1 Terrestrial Cloud

The following use case has been included in order to illustrate the concept of a terrestrial cloud and describe how the technology can be used in the tactical environment:

A mounted patrol has taken a number of suspected insurgents into custody. Facial images of each of the insurgents have been taken by the patrol using their handheld mobile devices. The images have since been uploaded to a cloudlet for processing. Due to the number of cloudlets in the VANET the warfighters are able to quickly identify eight of the ten individuals in custody as known insurgents. The patrol is just about to release the remaining two individuals when they are joined by another patrol. The cloudlet in the new LAV automatically joins the existing VANET and is assigned the task of attempting to identify the remaining two individuals. Since the new mounted patrol has just returned from base, its database has been recently updated. Within minutes the new cloudlet is able to positively identify both of the individuals as known insurgents, including one who has recently been deemed a person of interest in a terrorist attack.

5.5.2.2 Battlegroup Cloud

The following use case has been included in order to illustrate the concept of a battlegroup cloud and describe how the technology can be used in the tactical environment:

During a routine patrol the battlegroup intercepts encrypted communications from what they suspect are pirates targeting merchant vessels. Due to poor communications the battlegroup is unable to transfer the message back to command HQ for cryptanalysis. Instead, the battlegroup uses the entire computing resources of the battlegroup, interconnected through a VANET, to perform a brute force attack on the encrypted communications. In a relatively short time the communications are decrypted, verifying the nature and intentions of the pirates. The battlegroup is then able to intercept the pirates and put a stop to their actions. The brute force attack would have taken substantially longer without the combined cloud resources of the battlegroup. In all likelihood, the computing power available in a single vessel would have been insufficient to decrypt the encrypted communications in time.

5.5.2.3 Airborne Compute Cloud

The following use case has been included in order to illustrate the concept of an airborne compute cloud and describe how the technology can be used in the tactical environment:

UAVs have been sent out regularly to survey a remote area in the mountains that may be home to a group of insurgents. The UAVs are equipped with a variety of sensors, including high-definition video cameras. Given the sheer volume of data gathered by the UAVs, it is unrealistic to download the raw data to the FOB given the low-bandwidth wireless communications link. While the raw data can be transferred off of the UAV once it returns to base, there could be up to a two day delay between the data being collected and it being transferred off of the UAV.

The airborne compute cloud allows the raw data to be processed during flight. Not only can each UAV process its own raw data but it can distribute the processing to other UAVs that are part of the airborne compute cloud. Instead of having to transfer massive quantities of raw data down to the FOB for processing, the airborne compute cloud can merely send the processed data to the

FOB. Not only is the processed data significantly smaller, but it is actionable by warfighters thereby decreasing the lag between data collection and action.

5.6 Current Initiatives

Current initiatives have been divided into two types – those focusing on MCC and those focusing on VANETs.

5.6.1 MCC Initiatives

MCC initiatives include the following:

- Command and Control Applications for the Decisive Edge;
- Connecting Soldiers to Digital Applications (CSDA);
- Enabling Battlefield Decision-Making in the Tactical Cloud;
- Joint Battle Command-Platform (JBC-P);
- Cloud Computing at the Tactical Edge; and
- Civilian Research Initiatives.

5.6.1.1 Command and Control Applications for the Decisive Edge

The U.S. Army Communications Electronics Research, Development and Engineering Centre (CERDEC) issued a Broad Agency Announcement (BAA) in August 2012 entitled Command and Control Applications for the Decisive Edge. The BAA, which is valid through July 2017, is a competitive process for the selection of research and development proposals in the area tactical cloud computing, and specifically in the area of enabling warfighters to access Command and Control (C2) and intelligence services with a wide variety of military computers of variable link capacities located anywhere on the battlefield.

Included in the BAA is a topic, which is sponsored by the CERDEC Command, Power & Integration (CP&I) Directorate, called Command & Control Tactical Cloud Computing Environment. It seeks to enable warfighters on the forward edge of the battlefield to use cloud computing to access important situational awareness information using data radios, wearable computers, rugged laptop computers, and other rugged mobile computing devices.

5.6.1.2 Connecting Soldiers to Digital Applications (CSDA)

CSDA is a pilot program managed by the U.S. Army's Brigade Modernization Command (BMC) in Fort Bliss, Texas. The purpose of the program is to equip soldiers with smart phones and other handheld electronic devices. The program supports a wide variety of mobile devices, including

iPhones, Android-based devices, Touch Pros, Palm Treos, iPads, Kindles, etc. The intent is that soldiers will use these mobile devices in order to have access to email and other forms of communication. Eventually, the goal is to deploy a secure network on the battlefield that will allow warfighters to view real-time intelligence and video from air surveillance systems on their mobile device. In addition, mobile applications would also be capable of tracking both friendly and enemy soldiers on a dynamic map.

While the program is focussed on the use of mobile devices, it does have a number of cloud connections. The program is examining the use of clouds to store data collected by smart phones. In addition, the program is exploring the possibility of warfighters using their smart phones to access cloud-based data systems for pertinent information such as sniper zones, IED hotspots, most common attack times, alternate routes, etc.

5.6.1.3 Enabling Battlefield Decision-Making in the Tactical Cloud

This research, which is being conducted by the Army High Performance Computing Research Centre (AHPCRC), is looking at *the use of cloud computing to get closer to the reality of the Warfighter having the “right” information, at the “right” time, at the “right” place, and displayed in the “right” format.*²³ This research is examining the use of cloudlets to execute applications on behalf of warfighters in a timely manner. Specifically, this research will focus on 1) algorithmic support for the specification of tactical cloudlets that can service the computational needs of warfighters and 2) the development of performance-engineered solutions that have decreased resource demands and, thus, can be migrated to mobile platforms.

5.6.1.4 Joint Battle Command-Platform

JBC-P, which is a U.S. Army Program Executive Office (PEO) Command, Control, Communications – Tactical (C3T) program, allows soldiers in vehicles, aircraft and command posts to track friendly forces and exchange messages in order to reduce the likelihood of fratricide. JBC-P offloads processing to conventional computer hardware and software integrated into tactical vehicles and aircraft. The program also runs Nett Warrior²⁴ on networked handheld devices. This mission command system will allow team leaders to keep track of friendly forces. JBC-P moved into the production and deployment phase in the Summer of 2012.

5.6.1.5 Cloud Computing at the Tactical Edge

This research presents a reference architecture and prototype implementation for mobile devices that exploit cloudlets. However, since the research was sponsored by DoD, it focuses on the tactical environment in which networks are limited and bandwidth is limited and inconsistent. Both the initial prototype and the revised prototype implemented a face recognition application. The main challenge identified by the prototypes was the rapid delivery of large application overlays to cloudlets.²⁵ Additional information can be found in the paper *Cloud Computing at the Tactical Edge* [Reference 25].

²³ <http://ahpcrc.stanford.edu/?q=research/project/enabling-battlefield-decision-making-tactical-cloud>

²⁴ Nett Warrior is an integrated dismounted leader situational awareness (SA) system for use during combat operations of the United States Army. http://en.wikipedia.org/wiki/Nett_Warrior

²⁵ This is discussed further in Section 7 of this report.

5.6.1.6 Civilian Research Initiatives

According to DARPA scientists, commercial communication capabilities surpassed the military capability in the late 1990s.²⁶ Therefore, it is extremely important to monitor civilian research initiatives in order to determine potential candidates for military deployments. The challenge lies in adapting this civilian technology to a tactical environment replete with hostile forces and little to no communications infrastructure. There are a number of civilian research initiatives in the area of MCC. While not all of them employ MCC in the strictest sense, nor can any of them be deployed in tactical environments in their current form, they all made significant strides to further develop this capability so that it may eventually be deployed in the tactical environment. These civilian research initiatives include the following:

- Mobile Assistance Using Infrastructure (MAUI) – MAUI is a Microsoft research project to enable mobile devices to use Central Processing Unit (CPU)-intensive, and data-intensive applications such as facial recognition and language translation. It attempts to overcome the resource limitations of mobile devices through the use of nearby resource-rich cloudlets. Additional information on MAUI can be found in *MAUI: Making Smartphones Last Longer with Code Offload* [Reference 14];
- CloneCloud – CloneCloud is a system, consisting of an application partitioner and execution runtime, that automatically transforms mobile applications so that they can take advantage of the cloud. It accomplishes this by partitioning the application at runtime so that a specific thread is migrated from the mobile device to a clone in the cloud. The thread is re-integrated back to the mobile device upon completion of processing. Additional information on CloneCloud can be found in *CloneCloud: Elastic Execution between Mobile Device and Cloud* [Reference 15];
- Code Offload by Migrating Execution Transparently (COMET) – COMET is a prototype that attempts to augment mobile devices with systems available on the network. Specifically, it allows virtualized code to be offloaded, thereby improving processing speed an average of 2.88 times for the nine applications tested. Additional information on COMET can be found in *COMET: Code Offload by Migrating Execution Transparently* [Reference 16];
- Elijah – Elijah is a prototype system for Just-In-Time (JIT) provisioning of cloudlets by an associated mobile device. The system is capable of provisioning a cloudlet with a non-trivial VM image in ten seconds through the use of VM synthesis and a series of optimizations. The authors have also released their VM synthesis code for the rapid provisioning of a custom VM.²⁷ Additional information on Elijah can be found in *Just-in-Time Provisioning for Cyber Foraging* [Reference 17] and *A Reference Architecture for Mobile Code Offload in Hostile Environments* [Reference 29];
- Mobile Cloud Hybrid Architecture (MOCHA) – MOCHA, which is a prototype for a mobile-cloudlet-cloud architecture, serves as a platform for a face recognition application that utilizes mobile devices to capture facial images and off-loads the face recognition to

²⁶ <http://online.wsj.com/news/articles/SB10000872396390444772804577621950655761214>

²⁷ The code is available at <https://github.com/cmusatyalab/elijah-cloudlet>

cloudlets. Additional information on MOCHA can be found in *Cloud-Vision: Real-time Face Recognition Using a Mobile-Cloudlet-Cloud Acceleration Architecture* [Reference 26] and *Benefits of Utilizing and Edge Server (Cloudlet) in the MOCHA Architecture* [Reference 27];

- Odessa – Odessa is a runtime for offloading mobile interactive perception applications (e.g., face recognition, object and pose recognition, and gesture recognition) to nearby cloud servers. Additional information on Odessa can be found in *Odessa: Enabling Interactive Perception Applications on Mobile Devices* [Reference 18]; and
- Remote Processing Framework (RPF) – RPF is a framework for automatically migrating tasks from a portable computer over a wireless network to a server and returning the results. The intent of RPF is to save battery power on the portable device. Additional information on RPF can be found in *The Remote Processing Framework for Portable Computer Power Saving* [Reference 19].

5.6.2 VANET Initiatives

There is relatively little research and prototyping of cloudlets that leverage VANETs at this time. However, related initiatives include the following:

- VANET and Cloud Research; and
- UAV Research.

Note - Wireless Access in Vehicular Environments (WAVE)

WAVE, which is also known as 802.11p, is an extension of the 802.11 wireless standard specifically for Intelligent Transportation Systems (ITS) applications between vehicles and between vehicles and roadside infrastructure. One possible use case is the broadcast of accident information to surrounding vehicles within a 500m range. The target is to accomplish this within 500ms of the accident occurring. The intent is that 802.11p will provide substantial advances in vehicular connectivity, including speeds up to 54 Mbps. Although WAVE does not specifically address VANET clouds nothing in the standard seems to preclude its use in this context.

5.6.2.1 VANET and Cloud Research

Some VANET and cloud research has taken place over the last few years. *A Survey on Vehicular Cloud Computing* [Reference 28] provides a good overview of this research. Of particular interest are the various applications that have been identified as possibilities for the technology. These include an airport as a datacentre, parking lot data cloud, shopping mall data centre, dynamic traffic light management²⁸, self-regulated High Occupancy Vehicle (HOV) lanes²⁹,

²⁸ Dynamic traffic light management allows the vehicles participating in the traffic to help resolve congestion. An example of this would be a traffic jam occurring after a sporting event. Dynamic traffic light management would allow the vehicles involved in the traffic jam to dynamically manage the traffic lights in order to resolve the congestion.

²⁹ Self-regulated HOV lanes could be dynamically established in order to ease the effects of traffic jams.

managing evacuation³⁰, road safety message³¹, easing frequent congestion³², and managing parking facilities.³³

The first three applications listed basically boil down to leveraging the cloud computing capabilities in parked cars in a data centre capacity. *The combination of a massive amount of unutilized resources on board vehicles, such as internet connectivity, storage and computing power, can be rented or shared with various customers over the internet, similar to the usual cloud resources.*³⁴

5.6.2.2 UAV Research

There have been a few different research efforts that involve using UAV as part of an airborne compute cloud. These include the following:

- *Coding-Based System Primitives for Airborne Cloud Computing* [Reference 22] – This PhD dissertation is notable for three reasons. First and foremost, the author developed a set of best practices for operating a UAV wireless networking testbed. Second, the author designed and implemented FlowCode, which is a reliable link layer that provides improved ground-to-UAV bulk data transport. Lastly, the author designed CloudSense, which is a network switch prototype that can be used to manage and maintain high utilization of airborne cloud resources; and
- *Maximizing Throughput of UAV-Relaying Networks with the Load-Carry-and-Deliver Paradigm* [Reference 23] – This paper examined a Load-Carry-and-Deliver (LCAD) approach using UAV to relay messages between disconnected ground nodes. Specifically, the paper proposed a lightweight LCAD protocol and calculated its throughput in a number of scenarios.

5.7 Capability Assessment

*Awaiting discovery is an entirely new world in which mobile computing seamlessly augments the cognitive abilities of users using compute-intensive capabilities such as speech recognition, natural language processing, computer vision and graphics, machine learning, augmented reality, planning and decision-making.*³⁵

Cloud computing can potentially save energy BUT ONLY for the applications where the data to

³⁰ Vehicles participating in the VANET cloud would work in conjunction with the emergency rescue response office to appropriately route traffic. The evacuation would factor in the availability of resources, including food, water, shelter, and gasoline.

³¹ Vehicles participating in the VANET cloud would be capable of querying the sensors of other vehicles that are in relatively close proximity in order to ascertain potential road hazards ahead, road conditions, etc.

³² The VANET cloud would be able to determine where bottlenecks have occurred and plot appropriate routes to circumvent the problem.

³³ Vehicles participating in a VANET cloud would be able to manage real time data of available parking spaces and direct drivers to these spots.

³⁴ *A Survey on Vehicular Cloud Computing* [Reference 28]

³⁵ *The Case for VM-based Cloudlets in Mobile Computing* [Reference 20]

*be transferred from the device to the cloud is small enough for the network bandwidth be able to handle the request (i.e. using small amount of processing power and bandwidth), and the amount of computation is large enough that it better be done off of device.*³⁶

MCC using cloudlets has considerable potential as there are a wide-range of applications for the technology. The challenge will be replicating the performance and results that have been achieved in lab environments to the tactical environment. However, the main bottleneck for tactical networks is bandwidth. Effort has been made to decrease latency by using WiFi or short-range radio over single-hop networks. Consequently, applications that require considerable processing and data access but little bandwidth are ideal candidates for the technology. Possible candidate applications include advanced mapping, facial recognition, language translation, augmented reality, direction finding fusion centre, image recognition, mission planning, route calculation, injury diagnosis, etc. Offloading these applications to the cloudlet will greatly reduce the energy consumed on the mobile device by the processor, memory, and storage.

Interconnecting cloudlets through the use of a VANET provides a number of benefits. Not only would this architecture extend the range of the capability but it would multiply the processing and storage available. Whereas it may be difficult to implement in a civilian environment, interconnecting vehicular cloud resources using a VANET is somewhat easier to implement in a military environment. In a civilian environment users may be hesitant to allow other users to utilize their vehicle resources unless it stands to benefit them, whereas in a military environment, warfighters have decidedly less choice in the matter.

In terms of the maturity evaluation factor, the cloudlet scores a 2 since the technology is undergoing extensive research and prototyping. In terms of the current level of adoption evaluation factor, the cloudlet scores a 1 since the technology has extremely limited deployment within the tactical environment. In terms of the current impact evaluation factor, the cloudlet scores a 2 since the technology has had little impact on warfighter effectiveness. In terms of the potential impact evaluation factor, the cloudlet scores a 5 since the technology will likely have a profound impact on warfighter effectiveness. The capability assessment for the centralized cloud can be seen in Table 3.

Table 3 - The Cloudlet Capability Assessment

Evaluation Factors	Cloudlet
Maturity	2
Current Level of Adoption	1
Current Impact	2
Potential Impact	5

³⁶ *Cloud Computing for Mobile Users: Can Offloading Computation Save Energy?* [Reference 21]

6 Pico-Cloud

6.1 Overview

This section of the report will examine pico-clouds. Specifically, this section of the report will examine the following aspects of pico-clouds:

- Capability Deficiency;
- Architectural Overview;
- Tactical Cloud Discussion;
- Use Cases;
- Current Initiatives; and
- Capability Assessment.

6.2 Capability Deficiency

Mobile devices carried by warfighters can be used to collect sensor data. This data can be collected using sensors such as the GPS device, accelerometer, light sensor, microphone, thermometer, clock, compass, camera, etc. In many cases (e.g., low-bandwidth, unreliable communication link) it is not energy efficient to transfer large quantities of information from the mobile device. Consequently, this data is trapped on the individual device, usually in an unprocessed form, until such time as it can be transferred off of it without having to be concerned with depleting the energy resources of the device. The delay between data collection and processing can have a detrimental impact on mission success.

6.3 Architectural Overview

A pico-cloud, which is illustrated in Figure 13 and Figure 14, consists of cloud resources located on the mobile devices of each warfighter that are shared via a MANET (see note below) in order to provide enhanced computation and storage capabilities. The neighbouring mobile devices work in a collaborative fashion to perform distributed processing as part of a larger challenge. However, since the energy consumed during network transmission can frequently be orders of magnitude higher than that during local processing, it is more efficient to process data locally than to transfer it between devices.³⁷ Consequently, pico-clouds are viable for scenarios in which each mobile device collects data locally, processes it locally and then shares it with the collective in order to solve a larger problem.

³⁷ *Hyrax: Cloud Computing on Mobile Devices using MapReduce* [Reference 7]

Note – MANET

A MANET is defined as a *self-configuring infrastructureless network of mobile devices connected by wireless*.³⁸ A MANET differs from a VANET in two important respects: mobility and constrained devices. Since people are much less mobile than vehicles, a MANET is unlikely to experience the rate of change that one would find in a VANET. A MANET also has significantly less computational ability due to the fact that people carry constrained mobile devices whereas vehicles may be equipped with more powerful computing technology. In a traditional MANET, mobile devices are used to route data traffic, whether the traffic is related to its own use or not.




Figure 13 – Pico-Cloud (1 of 2)

³⁸ http://en.wikipedia.org/wiki/Mobile_ad_hoc_network




Figure 14 – Pico-Cloud (2 of 2)

6.4 Tactical Cloud Discussion

A pico-cloud provides the ability to process data at the tactical edge and act upon the results immediately. Since mobile devices are constrained in terms of processing and storage, they are poorly suited for performing complex calculations that involve significant amounts of processing. A pico-cloud facilitates the distribution of processing tasks amongst participating nodes. By participating in the MANET, nodes extend the range to encompass additional nodes. Furthermore, if a mobile device is tethered to a long distance radio and the radio fails, the mobile device participating in the MANET can simply tether to another device's radio.

6.5 Use Cases

The following use case has been included in order to illustrate the concept of a pico-cloud and describe how the technology can be used in the tactical environment:

A patrol has been tasked with finding some radioactive material that insurgents plan to combine with conventional explosives in order to make a dirty bomb. Each member of the patrol is equipped with a cloud-enabled mobile device that also serves as a Geiger counter. The members of the patrol spread out through the area where the insurgents are suspected to have hidden the radioactive material. Each member's mobile device takes readings that are processed locally on the pico-cloud. However, pertinent readings are shared with the pico-cloud, which uses the combined readings to individually adjust the bearings of the individual members of the patrol. In relatively little time the patrol is able to zero in on the radioactive material.

6.6 Current Initiatives

Pico-cloud initiatives include the following:

- Content-Based Mobile Edge Networking (CBMEN);
- Hyrax; and
- A Virtual Cloud Computing Provider for Mobile Devices.

6.6.1 Content-Based Mobile Edge Networking

CBMEN is a DARPA project to create a private ad-hoc data network for mobile devices in tactical environments. The intent is that warfighters will use the network to share data, including intelligence information and imagery, with other warfighters without having to return to camp to access a central server. The CBMEN software that is installed on each mobile device (on both smartphones and military radios) accomplishes this by automatically sending data to other mobile devices within range via WiFi, cellular and radio frequencies. The software basically converts each mobile device into a server that generates, maintains and distributes content. Even though cloud computing is not employed, the resulting MANET provides the equivalent of cloud storage across the participating mobile devices. Furthermore, mobile devices can join or exit the network as they come in and out of range. Phase 1 of the project, which proved the concept using Android phones and Army Riflemen Radios, was completed at Fort A.P. Hill. Phase 2 of the project, which hopes to develop ways to improve the efficiency of the data transfer and strengthen security, is currently underway.

The field tests proved the concept works and highlighted the potential benefits of real-time information sharing. At one point in the testing, two squads on foot patrol came within communication range of each other. One squad had information about a simulated person of interest that the other squad was seeking. The CBMEN software, working in the background on the troops' mobile devices, automatically transferred the information from the first squad to the other, without the second squad having to ask for it. As the second squad entered a building where the person of interest was, the squad used that information to immediately identify and apprehend its target.³⁹

6.6.2 Hyrax

Hyrax is a thesis project conducted at the School of Computer Science at Carnegie Mellon University. The thesis is entitled *Hyrax: Cloud Computing on Mobile Devices using MapReduce* [Reference 7]. The premise of the thesis is that it is feasible with today's mobile devices and networking to provide mobile cloud computing in order to support system-wide goals. The author attempted to accomplish this by running a Hadoop client in a Java VM on an Android smartphone and then using a smartphone cluster (ten smartphones) to execute MapReduce tasks. While the author was ultimately successful in getting it to work, he found that Hadoop is "fairly heavy-weight" for current smartphone platforms.

6.6.3 A Virtual Cloud Computing Provider for Mobile Devices

The researchers developed a framework for a virtual cloud computing platform using mobile phones. They then implemented the framework in a prototype. The prototype leveraged Hadoop for distributed processing and used Extensible Messaging and Presence Protocol (XMPP) for communications between devices. Additional information about this research can be found in *A Virtual Cloud Computing Provider for Mobile Devices* [Reference 24].

³⁹ <http://itsecuritypro.co.uk/morestories/us-military-tests-secure-p2p-technology-in-the-field/>

6.7 Capability Assessment

A pico-cloud is interesting in that it is one of the few technologies that is somewhat easier to implement in a military environment than in a civilian environment. In a civilian environment users may be hesitant to allow other users to utilize their mobile device resources unless it stands to benefit them. However, the concept is imperfect given the current state of the technology. Consider the following three points:

1. Personal mobile devices provide significantly less (1000 to 5000 times less) computational power than servers;
2. Computation-intensive tasks run on mobile devices consume large amounts of battery power (when a processor's clock speed doubles, the power consumption nearly octuples); and
3. The primary limiting factor for mobile devices are batteries, which have failed to keep pace with other aspects of mobile device technology.

In terms of the maturity evaluation factor, the pico-cloud scores a 1 since the technology is undergoing limited research and prototyping. In terms of the current level of adoption evaluation factor, the pico-cloud scores a 0 since the technology has not been deployed within the tactical environment. In terms of the current impact evaluation factor, the pico-cloud scores a 0 since the technology has had negligible impact on warfighter effectiveness. In terms of the potential impact evaluation factor, the pico-cloud scores a 2 since, without a major technology breakthrough, the technology will likely have minimal impact on warfighter effectiveness. The capability assessment for the pico-cloud can be seen in Table 4.

Table 4 – Pico-Cloud Capability Assessment

Evaluation Factors	Pico Clouds
Maturity	1
Current Level of Adoption	0
Current Impact	0
Potential Impact	2

7 Future Research

Neither centralized nor de-centralized clouds are conducive to further research. They are both relatively mature technologies that for one reason or another militaries have been somewhat slow to adopt. Similarly, pico-clouds are not conducive to further research. While MANETs and personal computing devices for warfighters promise to improve warfighter effectiveness and are being pursued aggressively in the research community, given the current state of technology the addition of pico-clouds to MANET devices would add little capability improvement at this time.

That leaves cloudlets. Based on the capability assessment, cloudlets are not a very mature technology, have limited adoption, but have considerable unrealized potential. Consequently, they make a good candidate for future research. In their paper *Gearing Resource-Poor Mobile Devices with Powerful Clouds: Architectures, Challenges, and Applications* [Reference 10], the authors divided MCC into two main types: computation offloading and capability extension. Computation offloading refers to the practice of offloading parts of resource-intensive tasks to the cloud in order to overcome resource constraints on mobile devices. Capability extending refers to the practice of extending the capabilities of mobile devices by performing tasks on the cloud that are not possible on mobile devices.

In addition, there are a number of approaches that can be used to augment the computing capabilities of mobile devices. The predominant approach, application offloading, offloads computationally intensive or data intensive applications to VMs running in the cloudlet. This can be accomplished in a number of ways. One is a technique called VM synthesis. In this case an application overlay is offloaded from the mobile device to the cloudlet. The application overlay represents the difference between a base VM, with only the operating system installed, and a VM with the application installed. Using VM synthesis would necessitate that the mobile device be equipped with as many overlays as required to complete a mission. Furthermore, these application overlays would need to be transferred over potentially constrained communication links to cloudlets. However, the advantage of this approach, which is effectively PaaS, is that cloudlets would only need to be deployed with base VMs. This is desirable given the remote location of the cloudlet and the lack of readily available technical support. An alternate approach, which is effectively Software-as-a-Service (SaaS), is to deploy the cloudlets complete with applications so that the mobile device merely has to transfer data to the cloudlet. This facilitates the transfer but at the expense of complicating deployment of the cloudlet. Effort would need to be made to ensure that the cloudlet's applications were compatible with the mobile devices.

The next step in terms of future research would be to perform a detailed examination of MCC using cloudlets, with specific emphasis on the two different types and the various techniques for performing computation offloading and capability extension. The research would highlight the more promising approaches, propose further research to validate or invalidate them, and identify the security implications of the various approaches. In addition, the research could look at extending the range, storage and processing using VANETs.

8 Conclusion & Recommendations

Modern warfare is almost entirely reliant on the availability of up-to-date information and situational awareness. Accurate information provides a competitive advantage over one's opponent and ultimately is critical to mission success. However, as one gets closer to the tactical edge, the availability and quality of the information degrades considerably. This phenomenon is sometimes referred to as the "fog of war". Specifically, warfighters at the tactical edge are forced to make snap decisions in a chaotic environment without sufficient information. This situation, which is compounded by the fact that communications back to command are over low bandwidth links and are often unreliable, has a detrimental effect on the likelihood of mission success.

Tactical cloud computing provides a means by which cloud computing can be extended within the tactical environment, in order to provide greater computational resources and improve collaboration and information sharing. By moving cloud computing closer to the tactical edge, it provides warfighters with the ability to both process data and access historical data. This report examined and assessed four tactical cloud architectures. The results of the capability assessment can be seen in Table 5. The four tactical cloud architectures are as follows:

- Centralized Cloud – A centralized cloud offers enormous potential in terms of cost savings, flexibility and the ability to process the big data associated with ISR. One has only to look at the widespread adoption of this technology within the private sector in order to appreciate its potential within the military. In terms of the tactical environment, a centralized cloud does not have a direct impact. However, in terms of an indirect impact it would facilitate the rapid development and deployment of improved applications. In addition, the parallel processing provided by a centralized cloud will facilitate advanced analytics on large data sets, thus providing warfighters with mission critical information;
- De-centralized Cloud – A de-centralized cloud offers significant benefit to remotely deployed warfighters. In terms of deployment, it can be deployed in a matter of weeks instead of months or even years. Once deployed, it ensures that actionable intelligence ends up in the hands of the warfighter quicker than otherwise would have been the case if forced to access a centralized cloud over long distance, potentially low-bandwidth, communication links;
- Cloudlet – MCC using cloudlets has considerable potential as there are a wide-range of applications for the technology. It has been determined that applications that require considerable processing and data access but little bandwidth are ideal candidates for the technology. Offloading these applications to the cloudlet will greatly reduce the energy consumed by the mobile device processor, memory, and storage. Furthermore, this capability can be extended and multiplied through the use of VANETs; and
- Pico-Cloud – A pico-cloud is one of the few technologies that is comparatively easier to implement in a military environment than in a civilian environment. However, given the current state of technology pico-clouds do not add value to MANET deployments in the tactical edge.

Based on the capability assessments it was determined that cloudlets represent the most viable candidate for further research. Specifically, this report recommended that a detailed examination of this tactical cloud architecture be performed in order to gain a greater understanding of the techniques used for performing computation offloading and capability extension.

Table 5 – Tactical Cloud Architectures Capability Assessment

Evaluation Factors	Centralized Cloud	De-centralized Cloud	Cloudlet	Pico Clouds
Maturity	5	4	2	1
Current Level of Adoption	3	2	1	0
Current Impact	2	3	2	0
Potential Impact	3	4	5	2

References

- [1] P. Mell & T. Grance, *The NIST Definition of Cloud Computing*, SP 800-145, NIST, September 2011. Available at: <http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf> (date of access: 30 Nov 2013);
- [2] *Security Guidance for Critical Areas of Focus in Cloud Computing V3.0*, Cloud Security Alliance, 2011. Available at: <https://cloudsecurityalliance.org/guidance/csaguide.v3.0.pdf> (date of access: 30 Nov 2013);
- [3] Cloud Computing Strategy, Department of Defense Chief Information Officer, Department of Defense, July 2012. Available at: <http://www.defense.gov/news/DoDCloudComputingStrategy.pdf> (date of access: 30 Nov 2013);
- [4] W. Jansen & T. Grance, *Guidelines on Security and Privacy in Public Cloud Computing*, SP 800-144, NIST, December 2011. Available at: <http://csrc.nist.gov/publications/nistpubs/800-144/SP800-144.pdf> (date of access: 30 Nov 2013);
- [5] *Data Analysis Challenges*, JSR-08-142, MITRE Corporation, December 2008. Available at: <http://www.fas.org/irp/agency/dod/jason/data.pdf> (date of access: 30 November 2013);
- [6] K. Ha et al., *The Impact of Mobile Multimedia Applications on Data Centre Consolidation*, School of Computer Science, Carnegie Mellon University, October 2012. Available at: <http://reports-archive.adm.cs.cmu.edu/anon/2012/CMU-CS-12-143.pdf> (date of access: 30 Nov 2013);
- [7] E. Marinelli, *Hyrax: Cloud Computing on Mobile Devices using MapReduce*, School of Computer Science, Carnegie Mellon University, 2009. Available at: www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA512601 (date of access: 30 Nov 2013);
- [8] N. Balasubramanian, A. Balasubramanian & A. Venkataramani, *Energy Consumption in Mobile Phones: A Measurement Study and Implications for Network Applications*, IMC'09, Available at: <http://ciir-publications.cs.umass.edu/getpdf.php?id=904> (date of access: 30 Nov 2013);
- [9] D. Kovachev, Y. Cao & R. Klamma, *Mobile Cloud Computing: A Comparison of Application Models*, RWTH Aachen University, 2010. Available at: <http://arxiv.org/pdf/1107.4940.pdf> (date of access: 30 Nov 2013);
- [10] F. Liu et al., *Gearing Resource-Poor Mobile Devices with Powerful Clouds: Architectures, Challenges, and Applications*, IEEE Wireless Communications, June 2013. Available at: <http://grid.hust.edu.cn/fmliu/iwc2013-mcc-fangmingliu.pdf> (date of access: 30 Nov 2013);
- [11] M. Satyanarayanan, *Pervasive Computing: Vision and Challenges*, IEEE Personal Communications, August 2001. Available at: <http://www.cs.cmu.edu/%7Eaura/docdir/ieeepcs01.pdf> (date of access: 30 Nov 2013);

- [12] K. Ha et al., Cloud Offload in Hostile Environments, School of Computer Science, Carnegie Mellon University, December 2011. Available at: <http://elijah.cs.cmu.edu/DOCS/CMU-CS-11-146.pdf> (date of access: 30 Nov 2013);
- [13] X. Luo, *The Cloud-Mobile Convergence Paradigm for Augmented Reality*, Qualcomm Inc., 2011. Available at: http://cdn.intechopen.com/pdfs/24824/InTech-The_cloud_mobile_convergence_paradigm_for_augmented_reality.pdf (date of access: 30 Nov 2013);
- [14] E. Cuervo et al., *MAUI: Making Smartphones Last Longer with Code Offload*, MobiSys'10, June 2010. Available at: <http://www.cs.duke.edu/~ecuervo/downloads/maui.pdf> (date of access: 30 Nov 2013);
- [15] B. Chun et al., CloneCloud: Elastic Execution between Mobile Device and Cloud, EuroSys'11, April 2011. Available at: <http://berkeley.intel-research.net/maniatis/publications/2011EuroSys-CloneCloud.pdf> (date of access: 30 Nov 2013);
- [16] M. Gordon et al., *COMET: Code Offload by Migrating Execution Transparently*, University of Michigan, 2012. Available at: <https://www.usenix.org/system/files/conference/osdi12/osdi12-final-11.pdf> (date of access: 30 Nov 2013);
- [17] K. Ha et al., Just-in-time Provisioning for Cyber Foraging, MobiSys'13, June 2013. Available at: <http://krha.kr/data/publications/vmsynthesis2013.pdf> (date of access: 30 Nov 2013);
- [18] M. Ra et al., *Odessa: Enabling Interactive Perception Applications on Mobile Devices*, MobiSys'11, June 2011. Available at: http://www.cs.columbia.edu/~lierranli/coms6998-10Spring2013/papers/odessa_mobisys2011.pdf (date of access: 30 Nov 2013);
- [19] A. Rudenko et al., *The Remote Processing Framework for Portable Computer Power Saving*, University of California, 1999. Available at: <ftp://ftp.cs.ucla.edu/pub/ficus/geoff/sac99.ps.gz> (date of access: 30 Nov 2013);
- [20] M. Satyanarayanan et al., *The Case for VM-based Cloudlets in Mobile Computing*, 2009. Available at: <http://research.microsoft.com/en-us/um/people/bahl/papers/pdf/cloudlets09.pdf> (date of access: 30 Nov 2013);
- [21] K. Kumar & L. Yung-Hsiang, *Cloud Computing for Mobile Users: Can Offloading Computation Save Energy?*, School of Electrical and Computer Engineering, Purdue University, April 2010. Available at: http://ieeexplore.ieee.org/xpl/login.jsp?tp=&arnumber=5445167&url=http%3A%2F%2Fieeexplore.ieee.org%2Fxpls%2Fabs_all.jsp%3Farnumber%3D5445167 (date of access: 30 Nov 2013);

- [22] C. Lin, Coding-Based System Primitives for Airborne Cloud Computing, PhD Thesis for Computer Science at Harvard University, October 2011. Available at: http://www.eecs.harvard.edu/~cklin/ckl_thesis.pdf (date of access: 30 Nov 2013);
- [23] C. Cheng et al., *Maximizing Throughput of UAV-Relaying Networks with the Load-Carry-and-Deliver Paradigm*, 2007 IEEE Wireless Communications & Networking Conference, 2007. Available at: <http://www.eecs.harvard.edu/~htk/publication/2007-wcnc-cheng-hsiao-kung-vlah.pdf> (date of access: 30 Nov 2013);
- [24] G. Huerta-Canepa & D. Lee, *A Virtual Cloud Computing Provider for Mobile Devices*, ACM Workshop on Mobile Cloud Computing & Services: Social Networks and Beyond, June 2010. Available at http://www.beknowledge.com/wp-content/uploads/2010/10/eccbcA-virtual-cloud-computing-provider-for-mobile-devices_2010_Proceedings-of-the-1st-ACM-Workshop-on-Mobile-Cloud-Computing-and-Services-Social-Networks-and-Beyond,-MCS-10,-Co-located-with.pdf (date of access: 30 Nov 2013);
- [25] S. Simanta et al., Cloud Computing at the Tactical Edge, Carnegie Mellon University, October 2012. Available at: http://resources.sei.cmu.edu/asset_files/TechnicalNote/2012_004_001_28146.pdf (date of access: 30 Nov 2013);
- [26] T. Soyata et al., Cloud-Vision: Real-time face Recognition Using a Mobile-Cloudlet-Cloud Acceleration Architecture, University of Rochester, 2011. Available at: <http://www.cs.rit.edu/~jmk/papers/cloud-vision.pdf> (date of access: 30 Nov 2013);
- [27] Z. Dou, *Benefits of Utilizing an Edge Server (Cloudlet) in the MOCHA Architecture*, University of Rochester, 2013. Available at: http://www.ece.rochester.edu/projects/wcng/papers/theses/duo_MSthesis.pdf (date of access: 30 Nov 2013);
- [28] M. Whaiduzzaman et al., *A Survey on Vehicular Cloud Computing*, Journal of Network and Computer Applications, August 2013. Available at: <http://www.cloudbus.org/papers/Vehicular-Cloud.pdf> (date of access: 30 Nov 2013);
- [29] S. Simanta et al., A Reference Architecture for Mobile Code Offload in Hostile Environments, Software Engineering Institute/School of Computer Science (Carnegie Mellon University), 2012. Available at: http://elijah.cs.cmu.edu/DOCS/cloudlet_hostile_MobiCASE2012_camera_ready.pdf (date of access: 30 Nov 2013);
- [30] Finch, L., Perrett, K., and Ross, V., *Protecting information in coalition cloud computing infrastructures*, Technical Report, TTCP-C3I-AG2-1-2013, December 2013;
- [31] *United States Government Accountability Office Report to Congressional Requesters - Agencies Need to Strengthen Oversight of Multibillion Dollar Investments in Operations and Maintenance*, General Accountability Office, November 2013. Available at: <http://www.gao.gov/assets/660/658794.pdf> (date of access: 31 December 2013);

- [32] A. Magar, *Data Protection in Multi-Tenant Cloud Environments*, DRDC Ottawa, March 2012;
- [33] D.J. Cappuccio, *Energy Savings via Virtualization: Green IT on a Budget*, Gartner, 10 November 2008; and
- [34] F. Dandashi et al., *Tactical Edge Characterization Framework – Volume 1: Common Vocabulary for Tactical Environments*, MITRE Technical Report, November 2007. Available at: http://www.mitre.org/sites/default/files/pdf/08_0037.pdf (date of access: 19 February 2014).

Annex A Cloud Computing Primer ⁴⁰

Cloud computing is defined by the National Institute of Standards and Technologies (NIST) as *a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.*⁴¹

This definition of cloud computing will serve as the basis for further discussion throughout this Annex. A good understanding of cloud computing concepts, as illustrated in Figure 15, is a prerequisite to fully appreciating the tactical cloud architectures discussed in this report. This section will examine the following aspects of cloud computing:

- Essential Characteristics;
- Service Models;
- Deployment Models;
- Key Concepts;
- The Case for Cloud Computing; and
- Disadvantages of Cloud Computing.

⁴⁰ Much of the information found in this Annex is based on information originally presented by the author in *Data Protection in Multi-Tenant Cloud Environments* [Reference 32]. However, the concepts addressed in this Annex have been significantly expanded.

⁴¹ *The NIST Definition of Cloud Computing* [Reference 1].




Figure 15 - Cloud Computing Concepts⁴²

A.1 Essential Characteristics

According to the NIST model, cloud computing has the following five essential characteristics:

- On-demand Self-Service – On-demand self-service is basically self-service provisioning. A user can automatically provision the computing resources that they require without third-party involvement;
- Broad Network Access – Broad network access refers not only to the network but the mechanisms and platforms used to access computing resources as well. The computing resources are network accessible using standard mechanisms from a variety of client platforms;
- Rapid Elasticity – Rapid elasticity refers to the ability to provision and de-provision computing resources commensurate with demand. The user can provision as many computing resources as they require for as long as they require;
- Measured Service – Measured service refers to the fact that computing resources consumed by an organization are monitored, controlled, and reported. This allows the organization to control and optimize its use of computing resources; and

⁴² This figure is based on one found in *Security Guidance for Critical Areas of Focus in Cloud Computing* [Reference 2]. However, it has been modified for use in this report.

- Resource Pooling – Resource pooling refers to the sharing of computing resources amongst multiple users. The user is typically unaware of the location and identity of other users of the shared resources.

A.2 Service Models

The service model used for a particular cloud dictates the degree of control the organization has over the computing environment. This section will examine four service models: Infrastructure as a Service (IaaS), Platform as a Service (PaaS), Software as a Service (SaaS), and Data as a Service (DaaS). In a traditional IT infrastructure, the user is responsible for the entire stack from the network and hardware up to the application. Cloud computing service models offer less control than in a traditional IT infrastructure. Of the cloud computing service models, IaaS offers the most control, followed by PaaS, and then SaaS. This can be seen in Figure 16. It is worth noting that IaaS is the foundation of all cloud services, with PaaS building upon IaaS, and SaaS in turn building upon PaaS. DaaS, which was not included in the NIST cloud computing model, was proposed in the *DoD Cloud Computing Strategy* [Reference 3].




Figure 16 - Service Model Control⁴³

A.2.1 IaaS

The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).⁴⁴

⁴³ This figure was taken from *Protecting information in coalition cloud computing infrastructures* [Reference 30]

⁴⁴ The NIST Definition of Cloud Computing [Reference 1]

IaaS provides the computing infrastructure, along with storage and networking. These computing resources are typically abstracted so that consumers are provided with Virtual Machines (VMs), virtual data storage and virtual network components. Consumers manage this computing infrastructure through management Application Programming Interfaces (APIs). The provider is responsible for securing the underlying infrastructure and abstraction layers, while the consumer is responsible for the remainder of the stack. Google, IBM, VMware and Amazon.com all provide IaaS offerings.

A.2.2 PaaS

The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.⁴⁵

PaaS, which sits on top of the IaaS layer, provides the computing platform and solution stack. It is intended to facilitate the deployment of applications by concealing the costs and complexity of the underlying platform. Consumers can build and deliver applications on the platform using programming languages and tools that are supported by the stack. The provider is responsible for securing the platform, while the consumer is responsible for securing the applications developed and hosted on the platform. Amazon Elastic Computing Cloud (EC2), Force.com, Google App Engine and Microsoft Azure are examples of PaaS offerings.

A.2.3 SaaS

The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.⁴⁶

SaaS, which sits on top of PaaS, provides the entire user experience, including the content, its presentation, the application(s), and management capabilities. It is sometimes referred to as "on-demand software" because it hosts the software centrally where it is accessed by users, normally using a web browser, over the Internet. In SaaS offerings, the security controls provided by the provider are typically negotiated into the service contract. It is worth mentioning that many SaaS providers don't use VMs. Instead, these service providers leverage a single logical instance of an application that is capable of handling large numbers of tenants. Google Docs, Salesforce.com and Yahoo mail are examples of SaaS offerings.

⁴⁵ *The NIST Definition of Cloud Computing [Reference 1]*

⁴⁶ *The NIST Definition of Cloud Computing [Reference 1]*

A.2.4 DaaS

DaaS is based on the concept that the product, data in this case, can be provided on demand to the user regardless of geographic or organizational separation of provider and consumer. Additionally, the emergence of service-oriented architecture (SOA) has rendered the actual platform on which the data resides also irrelevant. This development has enabled the recent emergence of the relatively new concept of DaaS.⁴⁷

DaaS abstracts the entire stack including the application, thereby allowing users to access data independent of any application. DaaS, as defined in the *DoD Cloud Computing Strategy* [Reference 3], encompasses two primary activities. The first is the continued implementation of the DoD Data Strategy and deployment of standardized data interfaces that make DoD information visible and accessible to all authorized users. The second is the incorporation of emerging “big data” technologies and approaches to effectively manage rapidly increasing amounts of information and deliver new insights and actionable information.

A.3 Deployment Models

This section will examine the various cloud computing deployment models. These deployment models *broadly characterize the management and disposition of computational resources for delivery of services to consumers, as well as the differentiation between classes of consumers.*⁴⁸ Specifically, this section will examine the following deployment models:

- Private Cloud;
- Community Cloud;
- Public Cloud; and
- Hybrid Cloud.

A.3.1 Private Cloud

*The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.*⁴⁹

A private cloud is operated solely for, and used exclusively by, a single organization. However, it can be managed by the organization on premises or even managed by a third-party off premises. Private clouds negate a few of the primary advantages of cloud computing. These include server consolidation through multi-tenancy and shared infrastructure/operational costs. However, they do provide the consumer with a considerable degree of control. This is especially important for defence departments.

⁴⁷ http://en.wikipedia.org/wiki/Data_as_a_service

⁴⁸ *NIST Guidelines on Security and Privacy in Public Cloud Computing* [Reference 4]

⁴⁹ *The NIST Definition of Cloud Computing* [Reference 1]

Many cloud service providers offer private clouds as well. For example, Amazon Virtual Private Cloud (VPC) is an EC2 instance that is completely isolated from their public cloud offering. It is basically an off-premise, virtual private network that an organization can use to host a subset of their IT infrastructure. The organization has complete control over this virtual network environment. Many other vendors, including VMware, Citrix, IBM, Oracle and Red Hat, provide software that allows organizations to host their own private cloud offerings. Private clouds are most often used in the tactical environment, although a community cloud (Section A.3.2) may be a possibility for coalition operations.

A.3.2 Community Cloud

The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.⁵⁰

A community cloud is a cloud infrastructure that is shared amongst a community of organizations with a common purpose. This common purpose can include a mission, security requirements, policy, or compliance considerations. The community cloud can be managed by the organizations or by a third party and may be located on-premise or off-premise. The infrastructure costs associated with a community cloud are borne by the entire community rather than a single organization. The Google government cloud and the Amazon Web Services (AWS) GovCloud are examples of community clouds. The Google government cloud is open to U.S. federal, state, and local government agencies. The AWS GovCloud is intended to allow U.S. government agencies and contractors to move sensitive workloads into the cloud by addressing their specific regulatory and compliance requirements.

A.3.3 Public Cloud

The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.⁵¹

A public cloud is operated for, and used by, a large group or even the general public. It is typically owned by an organization in the business of selling cloud services. This entity is commonly referred to as a Cloud Service Provider (CSP). Amazon EC2, Google App Engine, Salesforce.com and Windows Azure Services Platform are examples of CSPs.

A.3.4 Hybrid Cloud

The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or

⁵⁰ The NIST Definition of Cloud Computing [Reference 1]

⁵¹ The NIST Definition of Cloud Computing [Reference 1]

*proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).*⁵²

A hybrid cloud is a cloud infrastructure comprised of two or more clouds (private, community, or public) that are interconnected. This interconnectivity typically permits data and application portability between clouds. For example, an organization may use a private cloud to host a subset of their Information Technology (IT) resources, but integrate it with a security vendor's public cloud providing threat intelligence (e.g., Trend Micro's Smart Protection Network, Cisco IronPort SenderBase Security Network).

A.4 Key Concepts

This section will examine two concepts of importance to cloud computing. They are abstraction of resources and multi-tenancy.

A.4.1 Abstraction of Resources

Virtualization provides a means of implementing cloud computing through the abstraction of the underlying hardware resources. Consequently, virtualization is considered by many to be the foundation for cloud computing. However, as with multi-tenancy (Section A.4.2), it has not been identified as an essential characteristic of cloud computing by NIST. Virtualization, which introduces an abstraction layer between a physical resource and the service requesting the resource, allows multiple users to share the underlying physical resource. In addition, virtualization provides a degree of isolation so that users cannot interfere with each other's use of the physical resource. Virtualization is used to abstract common cloud computing resources, including compute, network, storage, and security resources, from the underlying hardware.

Figure 17, which is an instantiation of the cloud computing concepts, illustrates these virtualization concepts. The five layers below the abstraction (virtualization) layer, denoted by a solid line, constitute the physical hardware comprising the shared cloud computing resources. The four layers above the abstraction (virtualization) layer, denoted by a dotted line, represent primarily logical resources that have been assigned or allotted to tenants of the cloud. In addition, each tenant will have its own data that will be retained within the cloud. Tenant is a generic term for a user who is utilizing resources in the cloud either on a temporary or permanent basis. This section will examine each of the following layers in additional detail:

- 1) Compute - Compute resources refer to both the underlying physical servers and the tenant virtualized systems that they host. The compute layer below the abstraction layer, denoted by a solid line, represents the physical servers. The compute layer above the abstraction layer, denoted by a dotted line, represents the tenant virtualized systems or Virtual Machines (VMs). A VM encapsulates the logical representation of an information system, including hardware, operating system, and applications, into a file or set of files. The operating system and applications, which are fully abstracted from the underlying hardware by the hypervisor, interact with the hypervisor to access the physical hardware of the system;

⁵² The NIST Definition of Cloud Computing [Reference 1]

- 2) Network - The network layer below the abstraction layer, denoted by a solid line, represents the physical network infrastructure. The network layer above the abstraction layer, denoted by a dotted line, represents the tenant networks. While these networks can be physical, in all likelihood they will be logical networks (e.g., Virtual Local Area Networks (VLANs));
- 3) Storage - The storage layer below the abstraction layer, denoted by a solid line, represents the physical disk arrays and the Storage Area Network (SAN). The disk arrays, which are managed using specialized software and hardware controllers, typically consist of physical disks pooled together. The SAN is simply the network fabric that provides the connectivity from the physical hosts to the disk arrays. The storage layer above the abstraction layer, denoted by a dotted line, represents the tenant's allocated storage within the cloud. While this can be comprised of physical disk arrays and SAN, in all likelihood it will be compromised of logical partitions of a disk array and Virtual Storage Area Networks (VSANs). VSANs were invented as a means to partition a physical SAN into a number of logical components;
- 4) Security - The security layer below the abstraction layer, denoted by a solid line, represents the Physical Security Appliances (PSAs). PSAs are hardened physical systems that are used to perform a subset of security functions (e.g., firewall). The security layer above the abstraction layer, denoted by a dotted line, represents the security appliances assigned to that particular tenant. While these appliances may be physical in nature, in all likelihood they will be Virtual Security Appliances (VSAs). This practice is referred to as security virtualization. It abstracts the security layer from the underlying physical hardware by replacing PSAs with VSAs. A VSA is a virtual appliance ideally consisting of a hardened operating system and a single security application. VSAs typically consist of a single security application in order to be consistent with the principle of isolating security functions from one another; and
- 5) Management - This layer is comprised of the infrastructure required to manage tenant resources.




Figure 17 – Virtualization Concepts

A.4.2 Multi-Tenancy

Multi-tenancy is another characteristic not identified as an essential characteristic in the NIST cloud computing model. However, it has been identified as an important element of cloud computing by the Cloud Security Alliance (CSA). Multi-tenancy in cloud environments refers to the use of the same set of resources by multiple consumers, typically, although not necessarily, from different organizations. Multi-tenant cloud computing environments provide significant cost advantages over traditional computing environments due primarily to economies of scale.

A.5 The Case for Cloud Computing

Gartner predicts that by 2016 cloud computing will comprise the majority of new IT spending.⁵³

Now that the reader has a solid understanding of cloud computing it is important to understand what cloud computing provides over and above traditional computing technology. Consequently, this section will attempt to make the case for cloud computing by highlighting a number of benefits of the technology. These advantages include the following:

- Server Consolidation - The average CPU utilization for servers is approximately 10%. Cloud computing offers the opportunity to consolidate a number of servers on a single

⁵³ <http://www.gartner.com/newsroom/id/2613015>

system and as a result increase server utilization rates to the 60-80% range. Server consolidation results in a number of benefits to the organization. These include reduced hardware costs, reduced power consumption, lower cooling costs and space savings, as well as lower hardware maintenance, procurement, and support costs over time. It has been estimated by Gartner that virtualization can reduce energy consumption by up to 82% and space requirements by 85%;⁵⁴

- Server Portability – Cloud computing enables server portability by allowing VMs to easily be moved between physical systems. VMs can be moved to another system in order to perform hardware maintenance. They can also be copied in order to setup test environments or training centers. Due to their portability, VMs enable organizations to have a more dynamic infrastructure that can readily adapt to changing circumstances. In most cases, server portability is predicated on the availability of a common hardware infrastructure with uniform virtualization software;
- Business Continuity – Cloud computing offers a number of advantages in terms of business continuity. It allows servers to be spread across multiple systems in order to minimize the effects of a hardware failure. Cloud computing also supports dynamic failover in that in the event of hardware failure the VM would automatically be relocated to a new system. In addition, since servers are encapsulated into a single file they can easily be replicated and restored on any system. This is advantageous in terms of minimizing downtime but also in terms of backing up these files and restoring them to an alternate location in the event of disaster recovery;
- Provisioning - In order to facilitate deployment and maintenance most organizations use standard desktop and server images. Cloud computing allows a standard image to be created and then easily provisioned across desktops and servers. If the image becomes corrupted over time, it can be restored instantly;
- Legacy Support - Most organizations have legacy applications that were developed to run on a specific operating system. In many cases, these legacy applications are unable to run on newer operating systems and necessitate the use of a legacy operating system. In the case of servers, organizations are forced to host the legacy operating system/application on antiquated hardware that becomes increasingly difficult to maintain. In the case of desktops, organizations are forced to provide a separate desktop for the legacy operating system/applications. Cloud computing solves both of these problems by enabling organizations to run legacy operating systems/applications in VMs on modern hardware;
- Non-Persistent Environments – Organizations spend an inordinate amount of time building test, training, development and demonstration environments that attempt to mimic the current production environment. Cloud computing enables organizations to replicate complex production environments in a matter of minutes so that they can be used in these non-persistent environments. This can be accomplished by merely copying production VMs and provisioning them for the non-persistent environment;
- Separation/Isolation/Containment - The intent of virtualization is to isolate each virtual environment so that actions performed inside the virtual environment do not detrimentally impact other virtual environments or the host system. This is accomplished through the use of a hypervisor that prevents guest operating systems from directly

⁵⁴ *Energy Savings via Virtualization: Green IT on a Budget* [Reference 33]

accessing system resources and thus impacting other virtual environments. The hypervisor ensures that each guest operating system has access to a subset of the actual system resources and that excessive use of system resources by any one virtual environment is impossible. Furthermore, many organizations have a requirement to run applications that may expose the organization to higher levels of risk on the same platform as other applications. Likewise, some organizations have a requirement to run applications at different levels of security on the same system. Cloud computing allows organizations to run riskier applications in a ‘sandbox’ that is effectively isolated from other applications running on the same system. Examples of this include a VM for evaluating malware in order to determine its impact on a system or a VM for browsing sites that may contain malicious code. In the latter example, the state of the VM can be reset after each use, thus ensuring that any potential malware acquired during the session is destroyed; and

- Scalability – Since the underlying hardware resources are abstracted from the software running on them, an organization can easily upscale or downscale the IT requirements for a particular application as required.

A.6 Disadvantages of Cloud Computing

Cloud computing is not without its disadvantages. Consequently, this section will attempt to outline some of the disadvantages of the technology. These disadvantages include the following:

- Performance Bottlenecks – Due to server consolidation virtualization typically increases server utilization rates from 10% to the 60-80% range. While this may not result in performance bottlenecks during typical usage, high usage periods can result in bottlenecks unless precautions are taken. Potential performance bottlenecks in virtualized environments include the following:
 - Bus Architecture – Given that many VMs may share a given Peripheral Component Interconnect (PCI) bus, there may be a performance bottleneck;
 - CPU – A nominal CPU utilization in the 60-80% range can quickly exceed system capacity during nightly backup or during peak usage (users logging on in the morning). This is especially true if backup agents use local compression or encryption. Care must be taken to ensure that backups are appropriately staggered;
 - Disk – In many organizations, VMs are stored on a SAN. Systems access the SAN through a fibre channel Host Bus Adapter (HBA). Depending on the number of fibre channel ports and the number of VMs hosted on the system, multiple VMs may be forced to share the same storage controller. In certain circumstances this can exceed the available bandwidth. The situation is compounded if VMs are being backed up to the same storage device;

- Network – Systems are limited by the number of physical network ports that they can support. These network ports then need to be shared across all of the VMs hosted on the system. This problem is compounded if a subset of these network interfaces are used specifically for backup purposes;
- Backup Complications - Virtualized infrastructures are highly dynamic. VMs can move from one system to another in the event of live migration or unscheduled failover. This VM mobility complicates backup and recovery operations, especially if it is unscheduled. Backup software must communicate with the virtualization management software in order to ensure that it can locate the VM in order to perform backup operations. Backup is further complicated if the VM is moved during the backup process. The backup software must be sufficiently robust to recover from this situation and resume the backup from the new location;
- Infrastructure Costs - Somewhat surprisingly, significant infrastructure costs are incurred when organizations transition to a virtualized infrastructure. Most organizations expect infrastructure costs to be reduced due to server consolidation. While this is true over time, there will be sizeable initial infrastructure costs due to the requirement for new servers and shared storage, not to mention virtualization software costs. In all likelihood organizations need to standardize on new servers for virtualization. These new servers are typically equipped with lots of memory, sufficient network interfaces and even hardware-assisted virtualization. Furthermore, shared storage is required to store all of the organization's VMs;
- VM Sprawl - Anyone who has used virtualization software is intimately familiar with the phenomena of VM sprawl. Due to the ease with which VMs can be created, organizations are quickly inundated with an inordinate number of VMs for production, demonstration, training and development. Unless the organization is properly equipped to handle these VMs they will quickly exceed the organization's ability to manage them;
- Management – One would think that by reducing the number of physical systems to be managed, one would reduce the overall management overhead for the organization. In the case of virtualization this is not entirely true. In fact, the opposite is true. Server consolidation through virtualization will reduce the number of physical systems to be managed but it does not reduce the number of logical resources to be managed. In addition, the use of virtualization also necessitates the use of hypervisors or host operating systems and virtualization management servers. Each of these new components needs to be managed. Consequently, effort must be taken to reduce the management burden by implementing specialized virtualization management servers and by integrating virtualized servers with enterprise management tools to include monitoring, reporting and basic configuration; and
- Security - Virtualized environments have many of the same security issues as traditional environments. However, virtualized environments have a number of security issues specific to virtualized environments. For example, the hypervisor introduces an additional layer that can potentially be compromised.

Annex B Wireless Communication Primer

This annex will examine a number of wireless communication concepts of importance to the reader. These include communication frequencies, mobile phone standards, and wireless networking standards.

B.1 Communication Frequencies

Communication devices use different frequencies to transfer data. As the frequency increases the range tends to decrease while the available bandwidth increases. Militaries around the world use all of the communication frequencies listed below. The most commonly used frequencies for communications are as follows:

- High Frequency (HF) – HF, which ranges from 3 megahertz (MHz) to 30 MHz, is used primarily for long distance communication (e.g., shortwave radio). The HF frequency is often used for military operations;
- Very high frequency (VHF) – VHF, which ranges from 30 MHz to 300 MHz, is used for long-range data communication up to several tens of kilometres (e.g., land mobile stations). The VHF frequency is often used for military aircraft communication, including both air-to-air and air-to-ground;
- Ultra-high frequency (UHF) – UHF, which ranges from 300 MHz and 3 gigahertz (GHz), propagates by line of sight. While it can be blocked by hills or large buildings, it is suitable for television broadcasting, mobile phones, walki-talkies and satellite communication. The UHF frequency is often used for military aircraft communication, including both air-to-air and air-to-ground. It is also used for military satellite communications; and
- Super-high frequency (SHF) – SHF, which ranges from 3 GHz and 30 GHz, is used for point-to-point communication and data links. These include wireless Local Area Networks (LANs), mobile phones, satellite communications, microwave radio relay links and short-range terrestrial data links. The SHF frequency is often used for military satellite communications.

B.2 Mobile Phone Standards

Mobile phones typically use UHF and SHF for communications. The primary mobile phone standards are as follows:

- 3G – 3G, which is the third generation of cellular technology, has a potential transfer speed up to 3 megabit per second (Mbps). 3G operates at the UHF frequency band between 700 MHz and 2.2 GHz; and

- 4G – 4G, which is the fourth generation of cellular technology, has a potential transfer speed up to 100 Mbps for high mobility communication (e.g., cars, trains) and up to 1 gigabit per second (Gbps) for low mobility communication (e.g., pedestrian). However, in practice transfer rates of 2 to 12 Mbps are more realistic. 4G operates primarily in the frequency bands between 700 MHz and 2.7 GHz.

B.3 Wireless Networking Standards

Wireless LANs typically use UHF and SHF for communications in the Industrial, Scientific and Medical (ISM) band, which is not licensed. The primary wireless LAN standards are as follows:

- 802.11 – 1 or 2 Mbps operating at 2.4 GHz;
- 802.11a – 1.5 to 54 Mbps operating at 5 GHz;
- 802.11b – maximum 11 Mbps operating at 2.4 GHz;
- 802.11g – maximum 54 Mbps operating at 2.4 GHz;
- 802.11n – 54 Mbps to 600 Mbps operating at 2.4 GHz or at 5GHz;
- 802.11ac – 1 Gbps operating at 5GHz; and
- 802.11ad – theoretical maximum throughput of up to 7Gbps over 60 GHz frequencies.

List of symbols/abbreviations/acronyms/initialisms

AHPCRC	Army High Performance Computing Research Centre
APC2	Army Private Cloud
API	Application Programming Interface
AR	Augmented Reality
AWS	Amazon Web Services
BAA	Broad Agency Announcement
BMC	Brigade Modernization Command
C2	Command and Control
C3T	Command, Control, Communication – Tactical
CBMEN	Content-Based Mobile Edge Networking
CERDEC	Communications Electronics Research, Development and Engineering Centre
CMC	Cloud-Mobile Convergence
COMET	Code Offload by Migrating Execution Transparently
CP&I	Command, Power & Integration
CPU	Central Processing Unit
CSA	Cloud Security Alliance
CSDA	Connecting Soldiers to Digital Applications
CSG	Carrier Strike Group
CSP	Cloud Service Provider
CVBG	Carrier Battle Group
DaaS	Data as a Service
DARPA	Defense Advanced Research Projects Agency
DCGS-A	Distributed Common Ground System – Army
DSC	DCGS Standard Cloud
DIL	Disconnected, Intermittent and Low-bandwidth
DISA	Defense Information Systems Agency
DoD	Department of Defense
EC2	Elastic Computing Cloud
FCC	Federal Communication Commission
FMV	Full Motion Video
FOB	Forward Operating Base
GAO	General Accountability Office
Gbps	Gigabits per second
GHz	Gigahertz
GPS	Global Positioning System
HBA	Host Bus Adapter
HD	High Definition
HF	High Frequency
HOV	High Occupancy Vehicle
HQ	Headquarters
IA	Information Assurance
IaaS	Infrastructure as a Service
IED	Improvised Explosive Device
IP	Internet Protocol

ISM	Industrial, Scientific and Medical
ISR	Intelligence, Surveillance and Reconnaissance
ISSP	Integrated Soldier System Project
IT	Information Technology
ITS	Intelligent Transportation Systems
JBC-P	Joint Battle Command – Platform
JCD	Joint Capabilities Document
JIT	Just-In-Time
LAN	Local Area Network
LAV	Light Armoured Vehicle
LCAD	Load-Carry-and-Deliver
LDOS	Large Data Object
MANET	Mobile Ad-hoc Network
MAUI	Mobile Assistance Using Infrastructure
Mbps	Megabits per second
MCC	Mobile Cloud Computing
MHz	Megahertz
MOB	Main Operating Base
MOCHA	Mobile Cloud Hybrid Architecture
NB	Naval Base
NCOE	Network Centric Operating Environment
NIPRNet	Non-secure Internet Protocol Router Network
NIST	National Institute of Standards & Technologies
PaaS	Platform as a Service
Pbyte	Petabyte
PCI	Peripheral Component Interconnect
PEO	Program Executive Office
POD	Performance Optimized Datacentre
PRR	Personal Role Radio
RACE	Rapid Access Computing Environment
RPF	Remote Processing Framework
R&D	Research & Development
SaaS	Software as a Service
SAN	Storage Area Network
SEAD	Suppression of Enemy Air Defences
SEC	Software Engineering Center
SHF	Super High Frequency
SIPRNet	Secure Internet Protocol Router Network
SITREP	Situational Report
SOA	Service Oriented Architecture
TRANSTAC	spoken language communication and TRANSlation system for TACtical use
TS SCI	Top Secret Sensitive Compartmentalized Information
UAV	Unmanned Aerial Vehicle
UHF	Ultra High Frequency
V2I	Vehicle-to-Infrastructure
V2V	Vehicle-to-Vehicle
VANET	Vehicular Ad-hoc Network
VHF	Very High Frequency

VLAN	Virtual Local Area Network
VM	Virtual Machine
VPC	Virtual Private Cloud
VSAN	Virtual Storage Area Network
WAMI	Wide-Area Motion Imagery
WAVE	Wireless Access in Vehicular Environments
XMPP	eXtensible Messaging and Presence Protocol